MARKING IN A VISUAL OPERANT DISCRIMINATION IN PIGEONS

by

DAVID JOHN EDGAR

Submitted for the degree of
Doctor of Philosophy

Department of Psychology
University of Stirling

September, 1984

CONTENTS

ACKNOWLEDGEMENTS

It is a pleasure to thank Dr. D. A. Lieberman for his guidance, support and encouragement throughout this degree. I could not have wanted a better supervisor. I am also grateful to the technical staff of the Psychology Department and Animal House - C. Dow, J. Ison, B. Lavery, J. Nimmo, J. Russell, T. Stewart, B. Sutherland, P. Trollope - for all their assistance. The helpful and efficient approach to the typing of this thesis by Mrs. M. Abrahamson has been much appreciated.

I am greatly indebted to S. Hodgkinson for all the help and support she has given me. I would also like to mention Dr. G. Hall who first introduced me to learning theory and stimulated my interest in the area.

Finally, I would like to acknowledge the financial support I have received from the Science and Engineering Research Council.

ABSTRACT

Lieberman et al. (1979) and Thomas et al. (1983) found an enhancement of learning in a two-choice, delayed-reward, spatial discrimination task with rats if every choice response was followed by a salient event. They proposed that the salient event marked the preceding choice response in memory so that the subject was more likely to recall it when given reinforcement.

The experiments reported here tested whether a marking effect could also be found with a visual discrimination using different subjects - pigeons - and different apparatus - the operant chamber. Experiments 1 to 8 involved a discrete-trial procedure and a variety of parameters. Following choice responses with a marker did not facilitate learning. Two explanations for this outcome are proposed. One, early experience with the stimuli used as markers might have reduced their effectiveness. Two, the onset of the discriminative stimuli on the response key might have attracted the subjects' attention at a level such that marking provided no additional benefit.

Experiments 9 to 11 involved an invisible-trials, free-operant procedure. The discriminative stimuli remained available throughout each session and the onset of a trial was not indicated in any way. The first response following the start of a trial was designated the choice response. A clear marking effect was found, and the results could not be explained in terms of either arousal or generalisation decrement. This result indicates the generality of the marking phenomenon, and provides an automated procedure for its investigation.

Marking shows that an added stimulus during the gap between a response and reinforcement can sometimes facilitate learning. The implications of these results for models of information-processing in animals are discussed. Marking is examined in relationship to two areas where additional stimuli have interfered with learning - response overshadowing, and retroactive interference in delayed matching-to-sample.

TABLE OF FIGURES

CHAPTER 1

INTRODUCTION

1.1 THE MARKING EFFECT

Marking is a concept that was first introduced into the animal learning literature by Lieberman, McIntosh and Thomas (1979) to account for their findings in a series of delayed reward experiments. Lieberman et al. found that rats were able to learn a spatial discrimination in a maze if they were handled following their choice response. In the initial experiment a T-maze was used. Immediately after making a choice and entering one of the side areas, subjects were picked up and returned to the choice box or to their home cage. They remained there for 60 sec. At the end of the 60-sec delay, the subjects were picked up again and returned to the start box where they received reinforcement if their choice response had been correct. The rats showed a very high level of learning despite the delay of reinforcement.

In order to explain this result Lieberman et al. suggested a marking hypothesis that drew upon Kamin's work in the field of classical conditioning. Kamin (1968, 1969) proposed that for an increment in an associative connection to occur, the unconditional stimulus (US) had to provoke a subject into a backward scan of its memory store for recent stimuli. He suggested that this search allows the subject to identify the conditional stimulus (CS) that has a predictive or causal relationship to the US, and that this scan is prompted only by an unexpected, surprising US. The experiments reported by Kamin (1969) employed a standard blocking paradigm. In this paradigm a subject is first trained with a CS, A, alone. It is then trained to respond to a compound of A and a second CS, B, with the result that on testing, virtually no conditioning is found to have occurred to B. In one experiment conducted by Kamin (1969) a standard blocking paradigm was used but with a modification that made it rather similar to Lieberman et al's (1979) delayed discrimination task. Each compound trial was followed with two reinforcements - one immediately after the CS, the other 5 sec later. This second, unpredicted, US completely eliminated the blocking effect.

Lieberman et al. argued that a similar analysis could be applied in instrumental learning situations. They extended Kamin's analysis in two ways - first to include not only USs but "any salient and unexpected stimulus", and second, to allow the examination of earlier responses as well as stimuli in the backward memory scan. Thus in the maze experiment described earlier (Lieberman et al., 1979, Experiment 1), the experience of being picked up following a choice response was likely to be a highly salient event for the rat, either because of the stress involved in being handled or because of the reinforcing qualities that handling may have acquired. Such a salient event would initiate a search through the rat's memory store in an attempt to identify what was responsible for the animal being picked up. Of the various responses that could be identified in this search the most likely is the choice response that had immediately preceded it. The choice response could therefore be described as having been marked in the animal's memory because of the additional attention and rehearsal it received as a consequence of being handled.

The next stage in the analysis of the maze experiment is to consider what happens when the animal is rewarded at the end of the delay following a correct choice. Lieberman et al. (1979) suggest that, (as in Kamin's classical conditioning situations), reinforcement would also cause the rat to initiate a backward scan through its memory store. They argue that an event which is likely to be identified in such a search is the marked choice response. This view is based in part upon a generalisation from research on the factors that influence recall in other contexts. In particular they point to evidence from work on verbal memory in humans and the von Restorff effect which suggests that items which are particularly distinctive when presented are more likely to be recalled subsequently. Baddeley (1976, p.269) cites an unpublished experiment involving a free recall task in which the crucial item was printed in a different colour. This item was recalled significantly better than other items in the list, (see also Jones and Jones, 1942). Green (1956, 1958) reports experiments in which a numerical item appeared in a list of verbal items and vice versa. The unexpected, distinctive items were recalled better than the other items in the list. Green attributed this to the "surprising" nature of these items. Both Baddeley (1976) and Green (1956) suggest that the easier recall reflects the additional attention and greater rehearsal that the surprising event received when presented, (see also Jenkins and Postman, 1948). Lieberman et al. recognise that events other than the marked choice response may also be examined during a memory search initiated by reward. However, they argue, only the choice response will eventually become associated with reinforcement because its correlation with reinforcement is better than that of other events.

In summary then, Lieberman et al. (1979) propose that handling facilitates delayed reward learning through a marking effect. The salient event of being picked up causes the animal to scan its memory for causal responses. The choice response which immediately precedes being picked up is identified and thereby marked in the animal's memory. Later, when reward is given, another backward search is triggered and the marked choice response is identified as the causal factor.

A number of alternative accounts of the result could be offered and these are examined in Lieberman et al. (1979) and in a subsequent paper by Thomas, Lieberman, McIntosh and Ronaldson (1983). One possibility is that the results which were obtained were a consequence of experimenter bias. The expectations of the experimenter could have influenced the way in which subjects were handled. When learning was anticipated, subjects may have been handled more gently following correct choices than incorrect ones and this may have shaped their behaviour. Lieberman et al. (1979) expressed scepticism about the validity of this account for several reasons. Attempts were made to reduce variations in handling with the experimenter always following a stereotyped procedure for handling subjects. Also in their first experiment Lieberman et al. found learning in a group where they expected no learning. Furthermore, an experimenter-bias account of the results would have to argue that although expectations had no influence on the outcome of the first experiment, they did in the second and third experiments that Lieberman et al. report. However, in order to reduce the possibility that experimenter bias might have been responsible for their results, Lieberman et al. ran an experiment using other salient events as markers. This experiment (Lieberman et al., 1979, Experiment 4) involved a spatial discrimination task in a maze with a 2-min delay of reward. (A ground plan of the maze used in Lieberman et al., 1979, Experiments 3 and 4, is given in Figure 1.) Three markers were used in different groups - handling, a 2-sec light, and a 2-sec burst of white noise. The duration of both the light and

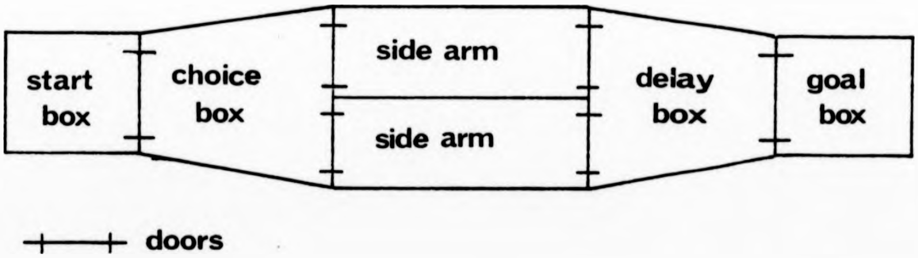

Figure 1

Ground plan of the maze used in Lieberman et al. (1979, Experiments 3 and 4) and Thomas et al. (1983).

the white-noise burst were controlled with an electronic timer. All the groups learned except the control group, and the white-noise and light groups performed better than the handling-marker group. Given that the duration and intensity of the marker in the light and white-noise groups were controlled automatically, there remained only one potential source of experimenter bias and that was in the timing of marker presentations. This possibility arose from the fact that the experimenter was responsible for initiating the marker by pressing a button. Therefore although the scope for invoking experimenter bias to explain the learning that occurred is very limited, it cannot be ruled out altogether. The demonstration of a marking effect with white noise and light in this experiment also provided support for Lieberman et al.'s (1979) view that any salient event could act as a marker, and not just handling. All three events seemed capable of initiating a backward scan through memory.

The marker acting as a secondary reinforcer is another possibility that Lieberman et al. (1979) examine. In their first experiment animals were picked up at the end of a delay and placed in the start box where reward was given if their choice response had been correct. Handling could thus have acquired secondary reinforcing qualities which were then present when the animal was picked up immediately following correct choice responses. However, as the animals were picked up after incorrect as well as correct choices, an account in terms of handling acting as a secondary reinforcer cannot explain differential strengthening of the correct response.

A third alternative to the marking interpretation focuses on the way in which marking may increase the subject's level of arousal. Evidence that increases in arousal level can improve learning in humans has been provided by work on paired-associate tasks by a number of researchers (see Berlyne, Borsa, Hamacher and Koenig, 1966; Hamilton, Hockey and Quinn, 1972; Kleinsmith and Kaplan, 1963, 1964; Walker and Tarte, 1963). Other studies with animals have also shown that increases in level of arousal as a result of drug administration facilitate learning (McGaugh, 1973). Berlyne (1967) reported two studies that showed an improvement in discrimination learning in animals if they were aroused by the injection of stimulants (Doty and Doty, 1966; Rensch and Rahmann, 1960), an effect that has also been found by Pare (1961). Therefore it may be that learning is facilitated in situations where the animal is more highly aroused and that contiguity between marking and the choice response is not crucial. Two experiments (Lieberman et al., 1979, Experiment 2; Thomas et al., 1983, Experiment 2) tested this alternative by running additional groups in which presentation of the marker was delayed until 30 sec after the choice response. If the arousal hypothesis is well-founded then learning in the delayed-marker group would be expected to be at least as good as, and possibly better than in the immediate-marker group. (Better performance might be anticipated as the delay in marker would probably result in the subject being more aroused at the time of reinforcement.) In fact the delayed-marker groups in both experiments showed significantly slower learning. Contiguity of marker and response is clearly necessary for the marking effect to occur.

Yet a fourth potential explanation that Lieberman et al. (1979) identify is based on the work of Muenzinger and his colleagues on the role of punishment. In a maze experiment with rats Muenzinger and Wood (1935) discovered that following choice responses with shock resulted in faster learning than in a control group. In a subsequent study, Muenzinger and Newcomb (1936) found a similar result when jumping a gap was substituted for shock. Muenzinger and Newcomb (1936) noted that in both the shock and jump situations subjects paused before making a choice, possibly because of the unpleasantness of these two events. They suggested that this extra exposure to the stimuli may have been responsible for the improvement in learning. (For a full discussion of Muenzinger's work see Fowler and Wischner, 1969.)

It could be argued that if being picked up is also an aversive experience, then the animals in Lieberman et al.'s marked groups may pause longer in front of the discriminative stimuli. The animals would then have greater exposure to the stimuli before making a choice response and may learn more rapidly as a result. However, Lieberman et al. recorded choice latencies for the different groups in each experiment - in no case was there any significant difference between marked and control groups. The view that greater exposure to the discriminative stimuli was responsible for the marking effect is therefore not supported.

In the experiments reported by Lieberman et al. (1979) the subjects received two markers on each trial - one immediately following a choice response, the other at the end of the delay. The use of two markers suggests a fifth alternative to the marking hypothesis. In explaining the basis for delayed-reward learning, a number of researchers have suggested an account based on memorial reinstatement (Cronin, 1980; Lett, 1979; Roberts, 1976; Spear 1973, 1978). Lett (1979) reported a series of T-maze experiments with rats. On making a choice response in these experiments, the rat enters a black or white chamber and information on the brightness of the chamber and on the response just peformed is stored in memory according to Lett. Later, when the rat is returned to the startbox at the end of the delay, memories of the chamber and of the choice response may be retrieved. If the animal is then fed in the startbox then the retrieved memories and the presence of reinforcement are represented at the same time and may become associated. Thus by a process of memorial reinstatement through retrieval cues the animal is able to learn the discrimination. A similar analysis is offered by Cronin (1980) to account for her findings with pigeons using a simultaneous visual discrimination task with delayed reward.

As well as evidence from animal work for the validity of a memorial reinstatement or retrieval cue hypothesis, there are also the results of human learning studies. Studies of memory for words have shown the role that retrieval can play in improving recall. One such study was conducted by Tulving and Osler (1968) who proposed an encoding specificity hypothesis to account for their results. They paired words with weak associates during training and found that presentation of the associate, the retrieval cue, at a later stage helped word recall.

Applying the memorial reinstatement and encoding specificity hypotheses to Lieberman et al.'s (1979) experiments suggests an alternative to marking. Rather than the first marker strengthening the memory of the preceding response it simply becomes associated with it. Then at the end of the delay period, when the marker is given for the second time its effect is to bring back into memory the first marker and, by association, the choice response. If food is given on that trial it also becomes associated with the response.

The retrieval cue analysis outlined above is clearly dependent upon a second marker and Thomas et al. (1983, Experiment 1) set out to test whether the second marker was indeed necessary. In the event both one-marker and two-marker groups learned equally well. Though it may be surprising that the second marker produced no additional learning, the essential result is that only one marker was required.

One final feature of the marking phenomenon which Thomas et al. identify can be discussed here. This concerns the question of whether or not markers facilitate memory for subsequent as well as preceding events. To test this possibility a maze experiment (Thomas et al., 1983, Experiment 2) was run in which one group received a 2-sec burst of noise immediately after the choice response and another group was marked immediately before making a choice. Both groups learned equally well.

One question that was posed by this result concerns how the marked-before group managed to learn the discrimination. If the marker initiated a backward scan through memory in this group, then behaviour that was irrelevant to solving the problem would be identified. In order to account for the learning in the marked-before group, Thomas et al. (1983, Experiment 2) drew upon work on the effect of arousal on attention (Easterbrook, 1959; Telegdy and Cohen, 1971). The impact of the marker in the marked-before group may have been to focus the animal's attention on the black and white stimuli in the two arms of the maze. These stimuli would then be more likely to be recalled when reward is given at the end of the delay.

A challenge to the backward scan account of learning in the marked-after group arose from the results of this experiment. As Thomas et al. point out, an attention analysis could be applied to the results for marked-after subjects as well as those marked before choices. The effect of the marker may not be to initiate a memory search but rather to alert the rat to the most obvious aspects of its environment. In the marked-after group, the subject will just have entered the side arm of the maze when it receives the marker, the effects of which may be to focus its attention on the black or white colour of the arm.

To test this attention explanation of learning in the marked-after group, Thomas et al. (1983, Experiment 3) ran an experiment in which there were no differential cues available in the side arms of the maze. Instead of one side arm being painted black and the other white, both were painted grey. The black and white stimuli were only available on the doors to the side arms. (Additional cues, such as the orientation of the maze within the room, were also controlled for.) If the attention account of learning was correct then animals should not have been able to learn the discrimination. The focusing of attention would not help the animal in any way to perform the task successfully. On the other hand, the marking hypothesis would still predict learning as a backward scan through memory should still identify the choice response. The results of the experiment supported the marking hypothesis.

The combined results of the two experiments just discussed led Thomas et al. (1983) to the conclusion that a salient event prompts two processes. The first of these is a backward scan through the memory store to identify causal or predictive cues. The second is a focusing of attention on external events. (It is perhaps interesting to note that in the experiment involving marked-before and marked-after groups, learning in the marked-after group - which could benefit from both processes - was not better than in the marked-before group.)

Several alternatives to the marking hypothesis have been discussed - experimenter bias, secondary reinforcement, arousal, stimulus exposure, and retrieval cues. None is supported by the experimental work that

Lieberman et al. (1979) and Thomas et al. (1983) report. An attention hypothesis has also been formulated which extends, but does not replace, the marking account of how a salient event promotes learning. It seems clear, (at least in the situations which have been discussed), that following a choice response with a salient event can facilitate learning by initiating a backward scan of memory to identify predictive cues.

The marking hypothesis has a number of potential implications and these are discussed by Lieberman et al. (1979) and Thomas et al. (1983). At one level the hypothesis offers waysof interpreting a range of other phenomena. Lieberman and his colleagues suggest that marking might help explain one, why response-contingent shock can sometimes assist learning (Muenzinger, 1934; Muenzinger and Wood, 1935); two, the phenomenon of "quasi-reinforcement" (Neuringer and Chung, 1967); three, feedback stimuli and secondary reinforcement (Skinner, 1938); and four, discriminated avoidance (D'Amato, Fazzaro and Etkin, 1968).

At a theoretical level, Lieberman et al. (1979) propose that marking may represent a fourth function to be added to the three functions traditionally assigned to a stimulus, those of eliciting, discriminating, and reinforcing.

A second implication of these results is that Lett's (1973, 1974, 1975) application of the principle of situational relevance to her findings of delayed-reward learning is not appropriate. The theory of situational relevance (Revusky, 1971) may well be correct but it does not account for the experimental results reported by Lett. (See Lieberman et al., 1979.)

A third area is the possible conflict between the marking effect and current models of information-processing in conditioning, including that proposed by Wagner (1976, 1978, 1981). Thomas et al. (1983) interpret Wagner's rehearsal model as predicting that a marker should interfere with learning rather than facilitate it.

In summary, the marking hypothesis proposed by Lieberman et al. (1979) and Thomas et al. (1983) may provide an explanation for several phenomena and it raises a number of interesting questions about the way in which animals process events.

1.2 RESEARCH FRAMEWORK

The starting point for the experimental work reported here was the desire to establish a new paradigm for the study of marking. This paradigm was then to be used to investigate a number of questions concerning the marking effect.

The initial aim was to develop an automated procedure for studying marking in pigeons using an operant chamber and a visual discrimination task. It was hoped thereby to provide strong evidence for the generality of the effect that Lieberman et al. (1979), and subsequently Thomas et al. (1983), had reported. The experiments they conducted explored some of the parameters of the phenomenon. Three different delays of reward were used - 30 sec, 1 min and 2 min - and two maze designs were employed. A range of markers had also been successfully tried - handling, light and white noise. Nonetheless, all the work conducted by Lieberman et al. and Thomas et al. involved rats, was done in mazes, and used spatial discrimination tasks.

In my experiments I wanted to test further the robustness and generality of the effect by using a different species (pigeons), a different type of apparatus (an operant chamber), and a different kind of discrimination task (visual).

In designing the experiments consideration was given to the similiarities and differences between the maze experiments described earlier (Lieberman et al., 1979; Thomas et al., 1983) and the operant chamber apparatus used in these experiments. One significant difference between the two situations is that in the maze the choice response can only be made once. Obviously other behaviour in the maze will subsequently occur but it is fairly different from the initial choice. In contrast, a pigeon in an operant chamber can continue to peck at the key after its choice response and, in effect, repeat behaviour that is identical to its initial response. There would therefore be the possibility in a spatial discrimination task in an operant chamber that pigeons could learn through adventitious reinforcement of responses which happened towards the end of the delay. All that would be required for this to occur would be a slight inertial tendency in the response patterns of the subjects. (By an inertial tendency is meant a tendency for the bird to repeat its most recent response.) Thus on trials where the subject was responding late in the delay, an inertial tendency would lead to chance pairings of responses to the correct stimulus with reinforcement, and responses to the incorrect stimulus with non-reinforcement. This process would lead subjects to develop a preference for the correct simulus.

Some evidence to support the view that subjects may show an inertial tendency is provided by Lea and Morgan (1972, p.131). Using rats in a two-lever chamber, they showed that subjects "chose a lever and then stayed with it for some time before shifting". This tendency to repeat the most recent response occurred despite experimental conditions that reinforced shifting between levers.

One answer to the problem of learning through adventitious reinforcement would be to remove the discriminative stimuli from the response keys before the end of the trial. (This account assumes that the correct and incorrect stimuli are projected onto different keys - the precise apparatus used in my work is described later.) However, in a spatial task if the bird continued to peck the darkened key then it would still be able to learn, (despite the removal of the stimuli), through pairings of responses to the correct key with food. Therefore spatial discrimination tasks might not be sensitive enough to uncover a marking effect. Both control and marked groups may benefit so much from adventitious reinforcement that ceiling effects would mask the contribution of the marker.

In a visual discrimination task it is easier to avoid the potential problem of adventitious learning. Clearly if the discriminative stimuli are present on the response keys until the end of the delay, then an inertial response tendency could still produce learning. However, their removal during the delay makes direct adventitious reinforcement impossible - an advantage that is not available when using a spatial task. Therefore the use of a visual discrimination problem in the experiments reported here not only brought the theoretical benefit of allowing the generality of the marking effect to be tested, it also had the significant practical advantage of making it easier to avoid the ceiling effects and loss of sensitivity that it was feared might hamper spatial discrimination work in an operant chamber.

A second major reason for attempting to develop the paradigm described earlier was the desire to have an automated procedure for studying marking. This was seen as having a number of advantages. First, it eliminated all possibility of experimenter bias in running experiments. The use of light and white-noise markers in Lieberman et al. (1979) and the mechanical control of their intensity and duration reduced enormously any room for the expression of bias. However, since the experimenter was responsible for pressing the button that operated the marker, there remained the admittedly limited possibility that slight differences in the timing of markers following choice responses could still provide differential information to the animals. Markers automatically initiated by the subject's choice response would remove any chance of bias in the administration of the marker and also substantially reduce the handling of subjects by the experimenter.

Automation would also bring the benefit of making the experiments less time-consuming for the experimenter as they would require a shorter period to run and would not need the constant attention that work with mazes demands. An automatic procedure would make it easier to increase the number of subjects and the number of trials involved in an experiment. Thus success in achieving an automatic procedure would bring a number of practical gains, would allow the elimination of any possibility of experimenter bias, and would help show the generality of the marking effect.

Two additional features of all the experiments reported in this thesis can also be mentioned here. First, the way in which the standard operant chamber was modified for this work and second, the use of only one marker on each trial.

On the front panel of the chamber was a row of three response keys. It was decided that to display the two discriminative stimuli on the two side keys might make the task too easy for the purposes of this research. If the discrimination was one that could be learned very rapidly, then performance in the control group might be so good as to obscure any effect of the marker in the marked group. Accordingly, the side keys were never used in this series of experiments and were always dark. The centre key was modified by splitting it vertically down the centre. Responses to each side of the key were recorded separately, and the discriminative stimuli were independently projected onto each half. Thus in the red-green visual discrimination task used, one half could be lit red, the other green and vice versa. It was hoped that this modification would produce a task that was neither so easy that the control group learned without any difficulty, nor so hard that no group solved the problem.

The other design feature that remained consistent throughout the research was the use of a single marker following immediately after a choice response. No additional marker was presented at the end of the delay interval as was the case in some of the experiments performed by Lieberman and his colleagues. Evidence of the redundancy of the second marker, (from the point of view of obtaining a marking effect), was provided by Thomas et al., (1983, Experiment 1) in an experiment that showed no difference in learning between a one-marker and a two-marker group. As the inclusion of a second marker allows the possibility of a non-marking interpretation, (see earlier discussion of the retrieval-cue hypothesis), it was dispensed with.

CHAPTER 2

DISCRETE-TRIAL I EXPERIMENTS

2.1 EXPERIMENT 1

Experiment 1 had two major objectives. The first was to establish in general terms whether the paradigm that had been selected would be appropriate for studying marking. The results of a pilot study had been encouraging, suggesting that the kind of delayed-reward visual discrimination task used in this experiment would not, for instance, be so difficult that no subjects could learn it.

The second objective was to show not only that the paradigm was appropriate, but actually demonstrate a marking effect. This would then provide a basis for work on the nature of marking.

A red-green visual discrimination task was used with a 15-sec delay of reward. The procedure, which is illustrated in Figure 2, was as follows. The discriminative stimuli were presented on the centre key and the first response initiated the trial. Two groups of birds were used and in both groups the choice response was followed by the offset of the non-chosen side of the key. The chosen side remained lit for 3 sec. The intention of this part of the procedure was to ensure some compatibility with the maze experiments discussed earlier, (Lieberman et al., 1979; Thomas et al., 1983). In these, after exposure to the discriminative stimuli at the choice point, the rat entered one of the side arms of the maze. Here it was exposed, (in some experiments for a second or two, in one for 30 sec), to the black or white stimuli associated with the arm of the maze that it had chosen. The rat then passed into a grey delay box for the remainder of the delay. (A ground plan of the maze was given in Figure 1.)

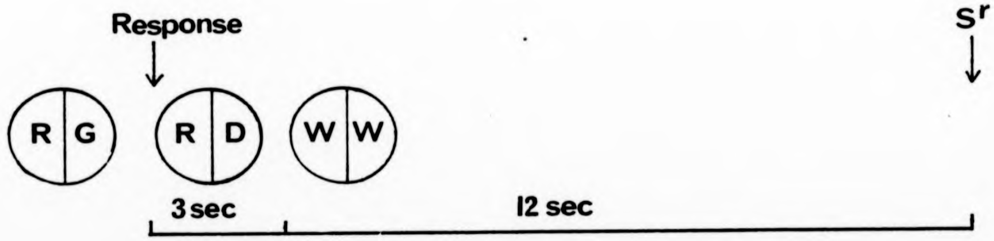

Figure 2

The design of Experiment 1. The circle represents the response key. The illumination of the key is indicated at each stage in a trial, (R = red, G = green, W = white, D = darkened). The trial illustrated is one in which a correct choice response is made to red. A marker also follows all choice responses in the marked group. (The same format is followed in subsequent figures showing designs.)

In the present experiment, after 3 sec had elapsed the colour on the chosen side of the pigeon's response key was removed in order to avoid the problem of adventitious reinforcement. Such reinforcement could have occurred had the subject continued to peck during the delay and if its response behaviour was subject to an inertial tendency. Chance pairings of pecks to the correct colour with food could have produced learning.

To prevent extinction of responding the entire key was illuminated with white light during the remaining 12 sec of the delay. Subjects were expected to continue pecking to the white key and thereby receive occasional pairings of a response with food reinforcement. Had the key been darkened during the delay then there was a danger that the birds' responding might have extinguished.

In the marked group only, the choice response was followed immediately by the offset of the houselight for 2 sec regardless of whether it was a correct or incorrect choice.

Method

Subjects

The subjects were eight adult Rock pigeons which were experimentally naive at the start of the experiment. They were housed individually in a different room from the experimental and control apparatus, and were maintained at 80% of their free-feeding weight throughout the experiment. Water and grit were continuously available in the home cages.

Apparatus

A Lehigh Valley Electronics operant chamber for testing pigeons was used. It measured 35 cm × 30 cm (floor dimensions) × 36 cm (wall height). The chamber contained a front panel with a row of three response keys (spaced 8.5 cm apart, centre to centre, and 26.5 cm above the floor). The centre key had a diameter of 3 cm, and the two side keys diameters of 2.5 cm. In this, (and subsequent), experiments only the centre key was used. The centre key had been modified by splitting it in half vertically down its centre. Responses to each half of the key were recorded independently. A multistimulus projector mounted behind the centre key was used to illuminate the key. Each half of the key could be illuminated independently with red, green, or white light. A red half-key emitted 170 μW of light, the green 110 μW and the white 110 μW.

In the centre of the front panel was a 6 × 5 cm food magazine whose lower edge was located 9.5 cm above floor level. The food magazine was illuminated when it made available wheat grain. The houselight was also mounted in the centre of the front panel 3 cm below the ceiling. The light, which was projected upwards, was provided by a 2.8-W bulb. It was continuously illuminated during each session except during the operation of the food magazine in the magazine training, autoshaping and pretraining stages of the experiment.

The chamber was placed in a sound- and light-attentuating shell with a ventilating fan behind, and to the left of, the front panel. Additional background noise was provided by white-noise from a speaker located behind the left-hand side of the front panel and 16 cm above floor level.

Most of the events of a day's session were controlled using Grason-Stadler electromechanical relay equipment located in a control room adjacent to the experimental room. During the training stage, a tape-reader was used to ensure that each subject experienced green on the left-hand side of the key and red on the right-hand side of the key, and vice versa equally often. The order of presentation of the stimulus pairs was randomised.

Procedure

Mazagine training and autoshaping. The birds received a 2-day magazine training sequence. On the first day the bird was placed in the operant chamber where the food magazine had been raised and piled high with grain. The only light available initially was that from the magazine light. After the bird had eaten for about 15 sec, the magazine was lowered. Thereafter magazine deliveries were progressively shortened until the bird was eating during a 3-sec magazine presentation. During this process the intervals between presentations of the magazine were progressively lengthened until they were between 10- and 15-sec long. The first session was ended when birds reached a criterion of nine approaches to the magazine out of ten consecutive presentations. On the second day of magazine training the bird was given 3-sec food presentations on a variable-time (VT) 60-sec schedule for 60 min. (On a variable-time schedule the length of the interval between times when reinforcement is presented varies around the mean value specified by the schedule, in this case 60 sec.)

On each of the next two days the birds were trained to peck the key using an autoshaping procedure. The entire response key was lit for 4-sec periods on a VT 60-sec schedule and each key illumination was followed immediately by 3 sec access to food. Fifty percent of the key illuminations were green and the remainder red. The order of green and red illuminations was random. Each autoshaping session lasted for 30 min.

Pretraining. The purposes of the pretraining stage were to develop a good rate of responding and to identify any colour preference that each bird might have. Fourteen daily sessions of pretraining were given to each bird. For the first ten sessions the birds received 30 trials a day. A trial consisted of the illumination of the entire response key either red or green for 60 sec with food reinforcement available on a VI 60-sec schedule.[1] Green and red trials were given equally often with their order randomised. Intertrial intervals lasted 60 sec during which the key was dark and no food reinforcement was available. Most of the birds developed a very strong preference for the right-hand side of the key, a preference which could have interfered substantially with training in the visual discrimination.

The pretraining procedure was therefore modified for the final four daily sessions. The VI schedule was replaced with concurrent variable-ratio (VR) schedules on each side of the key. (On a variable-ratio schedule the number of responses required before reinforcement is made available varies around the average value that the schedule specifies. For instance, on a VR 20 schedule, every 20th response on average is reinforced. The word "concurrent" indicates that more than one response is available to the subject and that separate schedules are in effect for each response. In this experiment two types of response were available - pecks to the left-hand side of the key, and pecks to the right-hand side.) The VR schedules in pretraining were manipulated in a successful attempt to modify the very strong side preferences and produce a response pattern more

---

[1] On a VI, or variable-interval, reinforcement availability varies around the mean value specified by the schedule.

evenly distributed across both sides. The right-hand preference may have developed because both the white-noise speaker and the fan providing ventilation were located on the left of the chamber. Both noises may have been aversive to the pigeons causing them to tend to the right. An alternative possibility is suggested by Werner and Wapner's (1952) sensory-tonic field theory of perception, which proposes that "any 'percept' is determined by an interaction between sensory activity (i.e. occurring in the sense organs) and tonic activity (i.e. proprioceptive feedback from muscles)" (Thomas, 1969, p.3). Evidence that this theory is applicable to pigeons as well as humans was reported by Thomas (1969). The impact of the auditory stimulation provided by the white noise and fan might have been to increase muscular tension on one side of the pigeon's body. The result of this might have been a change in the orientation of the bird away from the noises and towards the right-hand side of the key. To avoid the difficulties of aversiveness or sensory-tonic, the white-noise speaker and fan were switched off. The speaker was turned off after the first day of VR pretraining and the fan after the second day. Both remained off throughout the remainder of the experiment.

Delayed reward training. The subjects were divided into two groups of four birds each. The groups were matched for level of colour preference and response rate. The matching was done on the basis of the birds' performance on the last session of pretraining. The first group was called the marked group and the second the control group. If the bird has shown a colour preference on the last day of pretraining it was rewarded for choice responses to its non-preferred side. Birds that had shown no preference were rewarded in a way which ensured that, as far as possible, equal numbers of birds in each group were rewarded for red and for green.

Both groups received 24 daily sessions of delayed reward training on a red-green visual discrimination. Each session contained 30 trials and each trial began with the illumination of the response key, half red, half green. On 50 percent of the trials green was on the left-hand side of the key with red on the right-hand side, and vice versa on the remaining trials. The first peck to either side of the key on a trial was designated the choice response. This response initiated the trial and was followed immediately by the offset of the colour on the non-chosen side of the key. In the marked group only, the choice response was also followed by the offset of the houselight for two seconds. Three seconds after the choice response both sides of the key were illuminated with white light and remained so for a further 12 sec until the end of the trial. If the initial, choice, response had been to the correct colour, the subject then received 3 sec access to food. The end of the trial was followed by an interval of 10 sec, (during which the response key was dark), before the discriminative stimuli were presented again.

Results

Figure 3 shows the percentage of correct choice responses averaged over blocks of four sessions for the marked and control groups. Both groups show a substantial improvement in performance over the course of delayed reward training.

The data shown in Figure 3 were subjected to a two-way analysis of variance for groups and blocks of four sessions. There was a significant change in percentage correct over blocks, $F(5,30) = 13.6$, $p < .01$. The groups did not differ in the percentages of correct

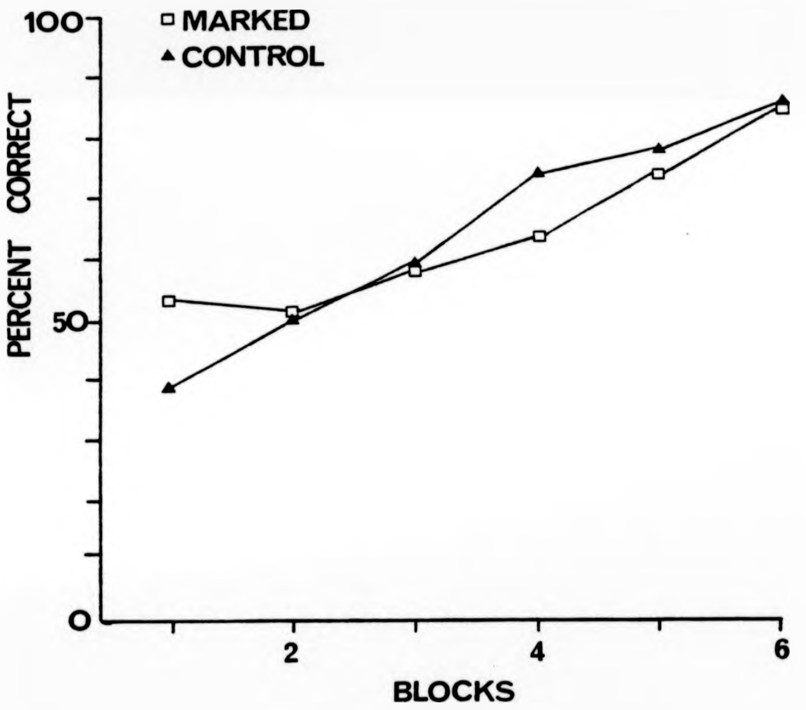

Figure 3

Mean percentage of correct choice responses over blocks of four sessions for each group in Experiment 1.

choices, $F(1, 6) = 0.01$, $p > .05$, nor was the Blocks X Group interaction significant, $F(5, 30) = 0.85$, $p > .05$.

Discussion

One aim of Experiment 1 was to establish whether a delayed-reward visual discrimination task with pigeons using a split response-key in an operant chamber might provide a suitable paradigm for investigating marking. The improvement in performance was substantial and the rate of learning was sufficiently slow to suggest that the design was sensitive enough to allow the effects of a successful marker to be revealed. Therefore in general terms the experiment offers support for the view that, with modification, this paradigm may well be suitable.

A second aim of the experiment was to provide a demonstration of the generality of the marking effect found by Lieberman et al. (1979) and Thomas et al. (1983). However, both the marked and control groups learned the discrimination and there was no difference between the groups in their rates of learning.

There are several potential explanations for this failure to obtain a marking effect. One possibility is that the discrimination task was too easy. Thus the control group was able to learn at a rate which prevented any benefit that the marked group may have derived from being revealed.

An alternative, or complementary, explanation focuses on the marker used in this experiment. In general terms it could be argued that the salience of a marker is important in determining how effective it is. The more salient a stimulus is, the greater might be the likelihood that it initiates a search through memory to identify possible predictive cues. The stimulus used as a marker in this experiment, houselight offset for 2 sec, simply may not have been salient enough. Intuitively this account does not seem very plausible and, in addition, there is evidence from work with delayed matching-to-sample tasks that pigeons are sensitive to changes in houselight illumination (Grant and Roberts, 1976). However, it remains possible that the failure to obtain a marking effect in Experiment 1 arose from using an insufficiently salient marker.

2.2 EXPERIMENT 2

In order to overcome the possible difficulty identified with the marker in Experiment 1, a second experiment was carried out with the same subjects. The aim was to see whether a more salient event would act as a more effective marker. Two groups of subjects were run and this time choice responses in the marked group were followed by a 2-sec burst of white noise in addition to houselight offset for 2 sec. Evidence of the capacity of white noise to act as a marker in other situations was provided by one of the experiments reported in Lieberman et al. (1979, Experiment 4). They found that a 2-sec burst of noise used as a marker in a spatial discrimination task with rats facilitated learning despite a 2-min delay of reinforcement. It was hoped that a combination of white noise and houselight offset in the present experiment would be an effective marker.

Method

Subjects and apparatus. The subjects and apparatus used in this experiment were those used in Experiment 1.

Procedure

Delayed reward training. The subjects remained in the same groups as in Experiment 1. The colour discrimination was reversed - thus subjects that had been rewarded for green choice responses in Experiment 1 were now rewarded for red choice responses, and vice-versa.

Subjects in the marked group had choice responses followed by a 2-sec burst of white-noise at 85 dB (SPL) and the offset of the houselight for 2 sec. In all other respects the training procedure in Experiment 2 was identical to that in the last experiment. Once again 24 daily sessions were given.

Results

The percentage of correct choice responses averaged over blocks of four sessions for the marked and control groups are shown in Figure 4. Both groups show a substantial improvement in performance over the course of delayed reward training.

A two-way analysis of variance for groups and blocks of four sessions was performed on the data shown in Figure 4. There was a significant change in percentage correct over blocks, $F(5, 30) = 14.29$, $p < .01$. The groups did not differ in the percentage of correct choices, $F(1, 6) = 0.20$, $p > .05$, and the Blocks X Group interaction was not significant, $F(5, 30) = 0.80$, $p > .05$.

Discussion

As in Experiment 1 both groups learned the discrimination and there was no demonstration of any marking effect. This failure may have occurred because attempts to produce a sufficiently salient marker were unsuccessful. An additional problem may have arisen from the fact that the subjects used in this experiment had already learned one visual discrimination, (involving the same colours), and were therefore already attending to the relevant dimensions of the task (Sutherland and Mackintosh, 1971).

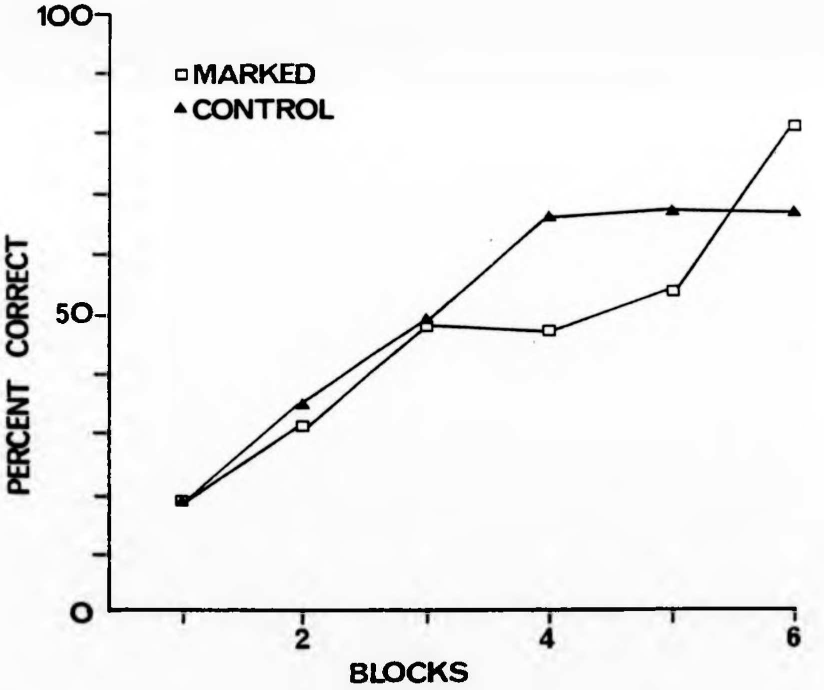

Figure 4

Mean percentage of correct choice responses over blocks of four sessions for each group in Experiment 2.

2.3 EXPERIMENT 3

For both groups in each of the experiments reported above, the initial choice response was followed immediately by the offset of the non-chosen colour on the other side of the response key. Therefore during the first three seconds of each trial the subject was exposed to one half of the key being blank, and the other half illuminated with the colour to which its choice response had been made. It may have been that this change in the state of the key was acting as a salient marker and hence producing equal learning in both groups. The intended markers - the offset of the houselight and the burst of white noise - may therefore have been redundant because a memory search would already have been initiated by the change in the state of the key. Therefore in Experiment 3 this source of non-differential marking was eliminated. In both the marked and control groups the first response on each trial no longer turned off the non-chosen side of the response key. Instead, both the red and green halves of the key remained illuminated.

A second potential source of non-differential marking in Experiments 1 and 2 occurred when the entire response key was illuminated with white light 3 sec after the initial choice response. (The key then remained white for 12 sec until the end of the delay.) If the birds continued to peck after their choice response and their responses were subject to an inertial tendency, then this change in key illumination may possibly provide a further marker. An inertial tendency would ensure that responses immediately preceding the change to white illumination tended to be the same as the initial choice response, thereby allowing adventitious marking in both groups. In order to reduce this risk of non-differential marking, the change was postponed in Experiment 3 until 5 sec after the trial began and the white light therefore only filled the final 10 sec of the delay.

In the marked group choices were followed by both a white-noise burst and the houselight offset, each for 2 sec. It was hoped that the presentation of these stimuli in combination from the outset of delayed reward training would represent an effective marker. If the marking hypothesis is correct, then a higher level or rate of learning would be expected in the marked group than in the control group. The procedure for both groups is illustrated in Figure 5.

Method

Subjects

The subjects were six adult Rock pigeons which were experimentally naive. They were housed individually in a different room from the experimental and control apparatus, and were maintained at 80% of their free-feeding weight throughout the experiment. Water and grit were continuously available in the home cages.

Apparatus

The apparatus was essentially the same as in Experiment 1 except for three changes in the operant chamber. One, the position of the loud speaker was moved and it was now mounted behind the centre of the front panel on the roof of the box housing the chamber. Two, the 4Ω speaker was used to supply white-noise markers. The third change was that the ventilating fan operated throughout the experiment and provided background noise of 72 dB (SPL).

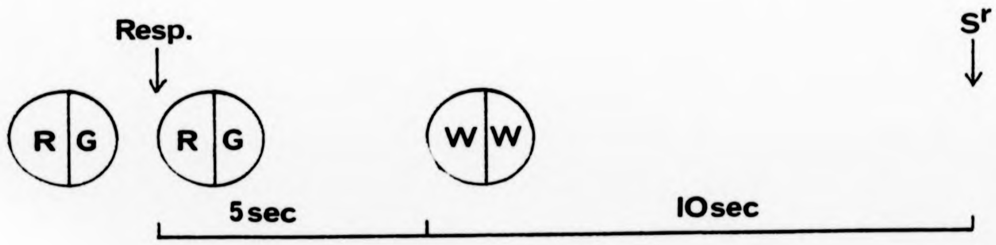

Figure 5

The design of Experiment 3. (For key see Figure 2.)

Procedure

Magazine training and autoshaping. As in Experiment 1 the birds received a 2-day magazine training sequence. The procedure on the first day was exactly the same as in Experiment 1, but on the second day free food was available on a VT 60-sec schedule for only 15 min. Each bird then received a single session of autoshaping lasting 30 min with the CS-US pairings presented on a VT 30-sec schedule.

Pretraining. Each pretraining session lasted about 30 min. During a session the bird received 44 trials lasting 30 sec each and the interval between trials was 10 sec. A trial consisted of one half of the key being lit with a colour, either green or red. Each of the four possible combinations of side and colour was presented 11 times and the order of presentation was randomly determined. During each trial reinforcement was available on a VI 30-sec schedule.

Four of the six subjects received a single session of pretraining; the remaining two required a further session before their responding was sufficiently well established to proceed to the training phase.

Delayed reward training. The subjects were divided into two groups of three birds each. The groups were matched for level of colour preference and response rate on the basis of performance in the final session of pretraining. During training subjects were rewarded for making choice responses to the colour they had preferred least in pretraining. If the bird had shown no colour preference, it was rewarded in a way that matched as far as possible the number of birds in each group that were rewarded for a particular colour. One group was called the marked group and the other, the control group.

Both groups received 32 daily sessions on the same red-green visual discrimination problem as in Experiments 1 and 2. The procedure was the same as in Experiment 1 except for the following changes. The number of trials in each session was increased from 30 to 50. The initial, choice, peck to the response key no longer turned off the illumination on the non-chosen side of the key. Instead, both sides of the key remained illuminated for a further 5 sec before changing to white for the final 10 sec of the trial. In the marked group only, the choice response was followed by a burst of white noise and houselight offset, both for 2 sec.

Results

The percentages of correct choice responses averaged over blocks of four sessions are shown in Figure 6. The marked group showed no improvement in performance over the course of the experiment, whereas the control group did improve somewhat.

The individual session data, (from which the blocked data shown in Figure 6 were derived), were subjected to a two-way analysis of variance for groups and sessions. There was no significant difference between the groups in overall correct responding, $F(1, 4) = 0.48$, $p > .05$, and no significant change in percentage correct over sessions, $F(31, 124) = 0.93$, $p > .05$. However, the interaction between experimental treatment and session was significant, $F(31, 124) = 1.73$, $p < .05$. Although the blocked data shown in Figure 6 suggest that this interaction may arise from a generally higher level of performance by the control group towards the end of the experiment, a simple main effects analysis of sessions reveals that control group performance is

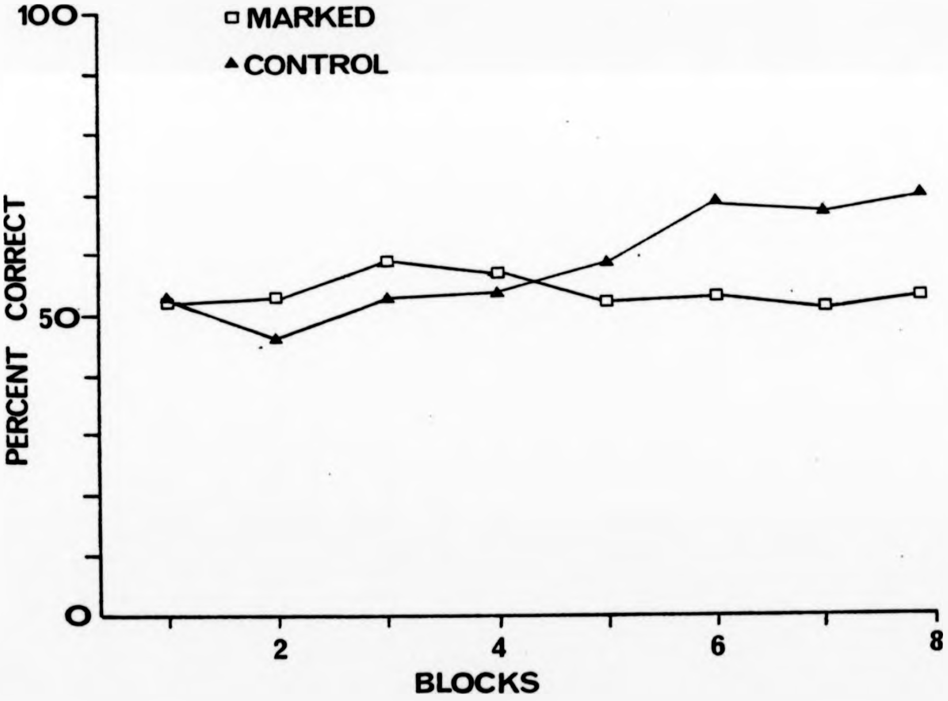

Figure 6

Mean percentage of correct choice responses over blocks of four sessions for each group in Experiment 3.

significantly higher only on Session 23, $F(1, 4) = 10.25$, $p < .05$, and Session 32, $F(1, 4) = 31.94$, $p < .01$.

Discussion

In Experiments 1 and 2 choice responses in both marked and control groups were followed by an event that may have acted as a non-differential source of marking. In Experiment 3 this potential problem was removed. The aim of the experiment was to test the hypothesis that following choice responses with a salient event in the marked group would promote better learning than in the control group. The results of the experiment do not support this hypothesis - in so far as a limited amount of learning did occur, it was found in the control rather than the marked group. This outcome poses a number of questions.

One issue raised is why the marked group failed to learn despite extensive training. It had been expected that a combination of a burst of white noise and the offset of houselight would represent a sufficiently powerful event to facilitate learning. One possible account of why this marker was not effective is suggested by consideration of general factors which might affect how well a marker operates. Salience has already been discussed as a factor, novelty could well be another. In his analysis of the classical conditioning situation Kamin (1969) pointed to the importance of a stimulus being surprising for it to initiate a memory search. Novelty may clearly affect how surprising a stimulus is and therefore the effectiveness of a marker may be determined in part by how novel it is. In the magazine training stage of this experiment, operation of the food magazine was paired with houselight offset. This experience could have influenced the operation of the houselight-offset element of the marker during delayed reward training. Birds may, for instance, have habituated to this stimulus especially as its duration during magazine training was similar to its length as a marker, (3 - as opposed to 2-sec).

This hypothesis - that both the salience and the novelty of a stimulus may be important in determining how likely it is to initiate a backward scan - could also account for the failure of the markers in Experiments 1 and 2 to produce additional learning. The subjects for these experiments had received about 100 food presentations and associated houselight offsets during magazine training. Also, in magazine training, autoshaping and most of pretraining white noise was used to mask background sounds. Habituation could therefore have occurred to both houselight off and white noise. Since in Experiment 1 the first of these stimuli was used as a marker, (and in Experiment 2 both were), habituation could have been responsible for the failure to obtain additional learning in the marked group.

Another account of the failure to get the expected level of marked learning in all the experiments reported so far also makes reference to the birds' experience during magazine training. Pairings of the offset of the houselight with food reinforcement at that stage may have endowed houselight offset with secondary reinforcing properties. If this had happened then it may have had a detrimental effect on the performance of the marked group during delayed reward training as the marker followed both correct and incorrect choices. (It should be noted that it could also be argued that if the marker had acquired secondary reinforcing properties, then its value as a marker should have been enhanced because the subjects would be more likely to attend to it and search for its cause.) However, if the houselight-offset marker had become a secondary reinforcer then one would expect a higher response rate in the marked groups (Mackintosh, 1974, p.233). A comparison of the two groups in Experiment 1 showed that in fact the opposite was true as the control group with a mean of 694 responses per session responded at a higher rate than the marked group with a mean of 536. The difference between the groups was not significant, $t(6) = 0.84$, $p > .40$. This evidence, then, does not provide support for the secondary reinforcer hypothesis.

It is possible therefore to identify three ways in which the marker may have been responsible for the failure to obtain learning in the marked group of Experiment 3. The marker may not have been sufficiently salient; habituation may have occurred to one of its elements; or, one element might possibly have become a secondary reinforcer. All of these possibilities are also relevant in analysing the failure to obtain marking in Experiments 1 and 2.

A second question that is posed by Experiment 3 is why the control group did somewhat better than the marked group. It may have been that the marker acted as a distractor, drawing the subjects' attention from the discriminative stimuli. However, this interpretation receives little support from Lieberman et al. (1979, Experiment 4) where a marking effect was found both with a 2-sec white-noise burst and with the switching on of a light for 2 sec. Another explanation is suggested by considering how the control group may have learned. One account available is that adventitious marking occurred at the point in the delay when the key lights changed from red and green to white. This account rests on the assumption of some inertial tendency in the birds' pattern of responding. It may have been that the effect of the marker following choices was to disrupt inertial tendency in the marked group and thus remove that basis for learning. A third possible explanation for the marked group's poor performance relative to the control is offered by Wagner's rehearsal model which was mentioned in Chapter 1 (Wagner, 1976, 1978; Wagner et al., 1973). However, the results of this experiment do not seem conclusive with regard to any of the potential explanations mentioned above. One problem with them all is that they fail to explain why the same markers in Experiments 1 and 2 did not similarly retard marked group learning relative to the control group. It would be useful to have an additional demonstration of the learning that was found in the control group of this experiment.

So far, two questions that were posed by the results of Experiment 3 have been discussed - why the marked group failed to learn and why the control group did better than the marked group. A third issue that is raised by the outcome of this experiment is why performance was so much worse than in the earlier experiments. Whereas in previous experiments the percentage of correct choice responses had risen by 40-50% in both groups, in this experiment only the control group improved and that by less than 20%. The possibility was raised earlier that the substantial level of learning found in Experiments 1 and 2 could be attributed to the offset of the non-chosen side of the key immediately after the choice response for all subjects. The evidence of the present experiment suggests that this may indeed have been acting as a marker and that its removal in Experiment 3 might explain the failure to achieve the same gain in performance.

An alternative, non-marking, account of the impact of switching off the non-chosen side of the response key is available. When the non-chosen side is darkened, the side that the bird made its choice response to remains illuminated for a further 3 sec. Evidence from a number of studies supports the view that the duration of stimulus exposure may be a factor that influences performance in discrimination tasks. Many investigations of matching-to-sample in pigeons have shown that matching accuracy is influenced by the duration of exposure to the sample stimulus (Carter and Werner, 1978; Grant, 1976; Maki and Leith, 1973; Maki and Leuin, 1973; Roberts and Grant, 1974). Therefore offset of the non-chosen side of the key may help learning not through marking, but through some kind of perceptual focusing as a result of additional exposure to the chosen colour. A final possibility is that both marking and perceptual focusing played a part in promoting learning.

2.4 EXPERIMENT 4

The results of the first three experiments pointed strongly to the importance for learning of the non-chosen side of the key being darkened after a choice response. Two interpretations of the role of this event were proposed, one of which was a marking account whilst the other relied upon non-marking processes. The marking version suggested that the change in key illumination acted as a salient event that initiated a memory search. The non-marking interpretation stressed the way in which additional exposure to the chosen colour could allow a beneficial focusing of the bird's attention.

The main aim of this experiment was to test these two alternatives. Three groups of pigeons were trained on the same red-green visual discrimination as in the previous experiments. The design of the experiment is illustrated in Figure 7. In the first, control group the choice response did not change the state of the key - both the red and green sides remained illuminated for a further 3 sec. In the second group - the half-key-marker-group - the choice response was immediately followed by the offset of the non-chosen colour. The chosen colour remained illuminated for a further 3 sec. The third group - called the whole-key-marker group - was one in which the choice response darkened both sides of the key for 3 sec. In all groups the key was illuminated white for the final 12 sec of the delay.

If the marking hypothesis was correct then the whole-key-marker group might be expected to show the most learning on the grounds that, intuitively, a complete darkening of the response key should act as a more salient event than its partial darkening. On the other hand, if the perceptual focusing through stimulus exposure account is accurate then the half-key-marker group should perform better than the whole-key-marker group.

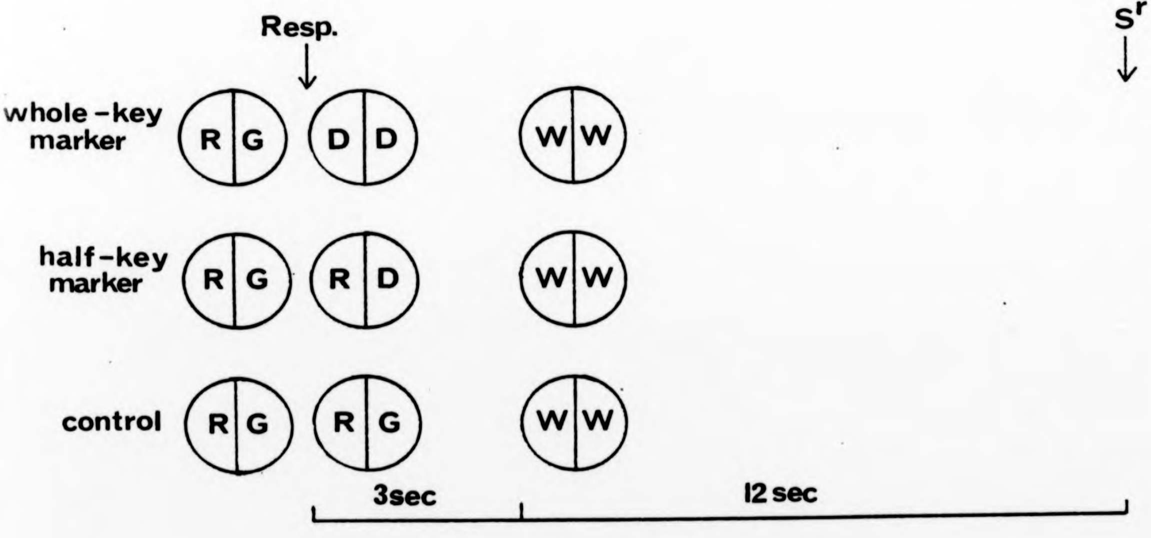

Figure 7

The design of Experiment 4.

The control group was included for two reasons. The first was that if both marked groups learned then one would want to be able to attribute their learning to a change in the state of the key. The second was to allow a partial replication of the control group in Experiment 3. The conditions in the two groups were not identical, (the change to a white key occurred at 5 sec in one, and 3 sec in the other), but were sufficiently similar to provide a test of the learning found in the last experiment's control group.

The inclusion of the half-key-marker group also provided a replication of the control group in Experiments 1 and 2.

Method

Subjects

The subjects were 21 adult Rock pigeons which were experimentally naive at the outset of the experiment. They were housed individually in a different room from the experimental and control apparatus, and were maintained at 80% of their free-feeding weight throughout the experiment. Water and grit were continuously available in the home cages. The subjects were run in two squads - the second squad of 12 birds was run after the first had finished their training.

Apparatus

For the first squad of birds the apparatus was the same as in Experiment 3 except that the ventilating fan was not used. In running the second squad of subjects an additional operant chamber was used. This chamber was identical to the first. It was controlled by BRS Digibit solid state equipment housed in the same room as the other control apparatus.

Procedure

Magazine training and autoshaping. Magazine training took two days. As in previous experiments, magazine training on the first day was controlled manually rather than automatically. On the second day, free food was made available on a VT 30-sec schedule for 20 min. Each bird then received a single session of autoshaping lasting 30 min with CS-US pairings presented on a VT 30-sec schedule. Four subjects required two autoshaping sessions.

During the magazine training stage presentation of food was no longer paired with offset of the houselight. Although houselight offset was not used as a marker in this experiment, there is evidence that changes in ambient illumination (the houselight) can be functionally equivalent to changes in illumination from a specific source (response key) in delayed matching-to-sample tasks (Grant and Roberts, 1976). Although Grant and Roberts were exploring an interference effect in a situation different from the visual discrimination used here, it is possible that a similar equivalence might still be found. As key illumination changes were to be used as markers in this experiment it was decided to avoid any potential problems by giving subjects no experience of houselight offset at this stage.

Pretraining. A single session of pretraining lasting about 30 min was given to each bird. The session contained 50 trials, each lasting for 30 sec, with an intertrial interval of 5 sec. During each trial both sides of the key were illuminated by the same colour and food reinforcement (3-sec access to wheat grain) was available on a VI 30-sec schedule. Red trials and green trials were presented equally often but in a random order.

Delayed reward training. All birds received 20 daily sessions of delayed reward training. Each session consisted of 50 trials and the interval between the end of one trial and the next presentation of the discriminative stimuli was 10 sec. Each trial lasted 15 sec from the time of the subject's first, or choice, peck to these stimuli.

Three groups of seven birds each were used. The groups were matched for level of colour preference, level of side preference, and response rate on the basis of the birds' performance in pretraining. In the second squad of birds, equal numbers of subjects from each group were run in each operant chamber. Where the bird had shown a colour preference in pretraining it was now reinforced for responding to the non-preferred colour. Birds that had not shown any preference were reinforced in a way that matched as far as possible the number of subjects in each group that were rewarded for each colour.

In the whole-key-marker group, the choice response darkened both sides of the response key for 3 sec. In the half-key-marker group, the choice response turned off only the non-chosen side of the key for 3 sec, and for the control group the choice response simply initiated the trial with the key remaining red and green for 3 sec. In all groups the key was lit entirely white for the final 12 sec of the delay, and correct choice responses resulted in 3-sec access to wheat grain at the end of the delay.

Results

The percentage of correct choice responses averaged over blocks of two sessions for all three groups are shown in Figure 8. It can

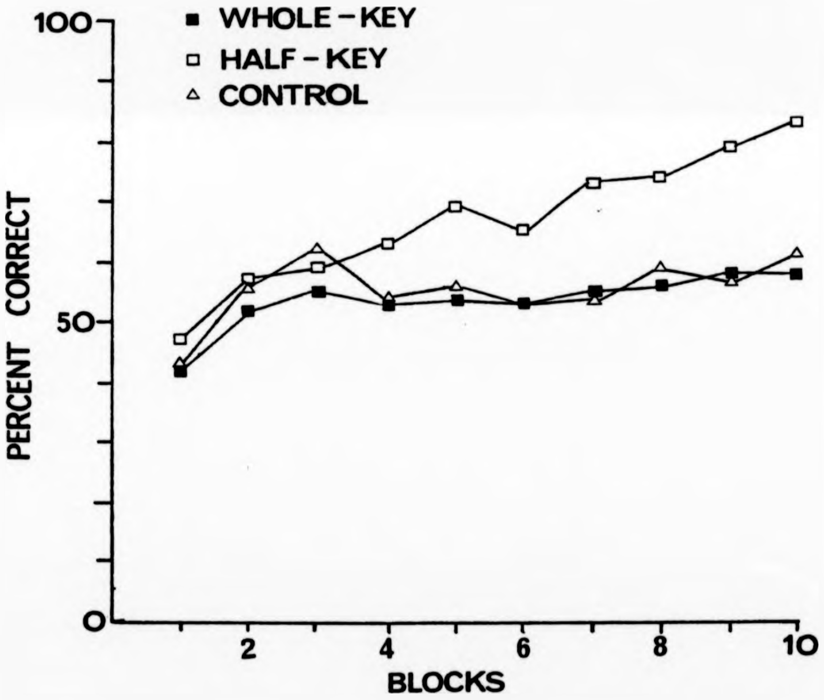

Figure 8

Mean percentage of correct choice responses over blocks of two sessions for each group in Experiment 4.

be seen that there is an improvement in performance for all groups over the course of delayed reward training. The level of improvement is clearly greatest for the half-key-marker group which increases from 45% correct initially to 85% correct on the final block. In contrast, both the whole-key-marker and control groups only improved from around 40% to 60% correct.

A two-way analysis of variance for the factors of experimental treatment and blocks of sessions was performed on the data shown in Figure 8. It confirmed that there was a significant change in percentage correct over blocks, $F(9, 162) = 6.28$, $p < .01$. The groups did not differ in their percentages of correct choices, $F(2, 18) = 1.36$, $p > .05$, and the Blocks X Group interaction was not significant, $F(18, 162) = 1.16$, $p > .05$.

In order to examine changes in performance over blocks a trend analysis was carried out. This showed that the linear, $F(1, 18) = 14.79$, $p < .01$, cubic, $F(1, 18) = 13.81$, $p < .01$, and quartic, $F(1, 18) = 6.27$, $p < .05$, trends were all significant. However, it is clear from Figure 8 that the overwhelming trend was linear and indeed the linear component of the trend accounted for 78% of the variation.

The linear trend of the half-key-marker group appears to be substantially greater than that of the whole-key-marker and control groups, although the difference does not quite reach significance, $F(2, 18) = 2.93$, $p = .079$.

Discussion

The main purpose of this experiment was to evaluate the marking and perceptual focusing accounts of the results of Experiments 1 to 3. These experiments had suggested a crucial role in learning for the effect of the offset of the non-chosen side of the key following a choice response. The outcome of Experiment 4 supports the perceptual focusing through stimulus exposure interpretation - had the marking account been correct then learning in the whole-key-marker group should have been superior to the half-key-marker group. In the event the whole-key-marker group was not only worse than the half-key-marker group, but also no better than the control group which experienced no stimulus change following choice responses.

A secondary aim of Experiment 4 was to attempt to replicate some of the earlier results of certain groups. In particular the half-key-marker group in this experiment was identical in design to the control group of Experiment 1. The results for the two groups were very similar - in the first experiment the control group improved from 40% correct to 85% correct over the course of training and a similar increase occurred in Experiment 4's half-key-marker group, up from 45% to 85%.

Another replication of interest was that of the control group in Experiment 3. The control group in the present experiment was slightly different in design, (the key changed to white after 3 as opposed to 5 sec), but the results were very close. In both groups performance at the end of training was about 20% higher than at the beginning. The success of both these replications is useful in providing confidence in the reality of the effects involved.

2.5 GENERAL SUMMARY AND DISCUSSION

The series of experiments reported so far was aimed at demonstrating the generality of the marking phenomenon and thereby finding a paradigm which would allow an exploration of the effect. The task employed in all the work described above was a red-green visual discrimination with a 15-sec delay of reward.

In the first two experiments a high level of learning was found but no marking effect was revealed - the performance in the marked and control groups was the same. This failure could have been because the task was too easy or the events used as markers - the offset of houselight and a burst of white noise - were not salient enough.

An alternative explanation was suggested in the third experiment. A potentially powerful source of non-differential marking was eliminated in that the choice response no longer darkened the non-chosen side of the key. Maximal differential marking was attempted by following choices in the marked group with both a burst of white noise and the offset of the houselight. The results were disappointing - not only was there no learning in the marked group, but its performance was somewhat worse than the control group's. Three accounts could be offered of why the marked group did not learn in the way that the marking hypothesis suggests it should have done. The marker may not have been salient enough, it may not have been sufficiently novel, and it could have acquired secondary reinforcing properties. (Incidentally, the difficulty that both groups had with learning the discrimination does not support the possibility raised by Experiments 1 and 2 that the task was simply too easy for marking to produce additional learning.)

The most striking aspect of the outcome of Experiment 3 was that learning was so much worse than in the first two experiments - a result that pointed strongly to the role of the offset of the non-chosen colour in promoting learning. One hypothesis about how this event may have influenced solution of the problem pointed to its potential as a marker. The other interpretation rested on a non-marking view that identified perceptual focusing arising from differential stimulus exposure as the learning mechanism.

These two accounts were tested in Experiment 4, the outcome of which was strongly in support of perceptual focusing. The group which experienced the greatest change in response key illumination after a choice performed no better than the control group and less well than the half-key-marker group.

The work reported so far leads to several conclusions. The first of these is that the major factor assisting learning in the experiments above was perceptual focusing, not marking. Events which one might have expected to function as markers on the basis of previous research (Lieberman et al., 1979; Thomas et al., 1983) did not do so here. A confounding factor in the first three experiments was the early experience that subjects had with stimuli which were subsequently used as markers or elements of markers. This experience may have affected the ability of these events to function as markers through habituation. (A second possibility - that the events acquired secondary reinforcing properties - can probably be dismissed given the failure to find a significant difference between the response rates of the groups in Experiment 3.) If it is the case that the early experience of the birds affected the properties of the markers, then the failure to get a marking effect rests upon relatively trivial factors.

The fundamental question that is raised is whether the absence of marking in these experiments is a result of technical difficulties of the kind described above, or whether it reflects a more profound limit to the generality of marking. Would an adjustment to the parameters of the task or the removal of confounding early experience eliminate the difficulty in getting marking, or does the problem rest with the species, kind of discrimination, or experimental situation being used?

CHAPTER 3

DISCRETE-TRIAL II EXPERIMENTS

3.1 EXPERIMENT 5

The learning that was found in Experiments 1 to 4 could not be attributed to marking. In some cases it was a result of perceptual focusing through differential stimulus exposure. In others it was either the same in both marked and control groups or superior in the control group and therefore attributable to some non-marking process. The aim of this experiment was to pursue an answer to the question of whether this failure to obtain a marking effect was due to the wrong choice of parameters or whether it reflected a more fundamental limitation on the phenomenon's generality.

Two changes were made to the design which it was hoped would allow a marking effect to be demonstrated. The first of these was a reduction in the delay of reward from 15 sec to 6 sec. The reason for this shortening of the delay was to make the task easier. The improvement in performance in earlier experiments, (except where perceptual focusing was possible), had been fairly limited with increases of only 20 percentage points over the course of training.

The second design change was the introduction of a different form of marker. In earlier work the markers used had been bursts of white noise and houselight offset, both for 2 sec, and the darkening of the entire response key for 3 sec. Although a 2-sec white-noise burst had been used successfully by Lieberman et al. (1979) in their maze experiments, it may not have been such an effective marker in the operant chamber work reported here. This is because in the maze experiment the rat was not able to repeat its choice response once it had entered a side arm, whereas in an operant situation the pigeon could continue to peck at the response key. Thus responses to both red and green stimuli may have been marked during a 2-sec burst of white noise and learning thereby made more difficult. A similar argument could be applied to the houselight-offset marker. A different problem with the latter marker may have been not that other responses were marked, but that a simple change from illumination to non-illumination was not sufficiently novel. The subjects had, after all, experienced changes in key illumination many times during the stages of the experiment before training and during training itself.

In order to overcome the possible problems of multiple response marking and insufficient salience, a new marker was used in this experiment. Both the houselight and the keylight were turned off for 0.5 sec following a choice response in the marked group, thus providing a very brief, complete darkening of the chamber.

Some support for the view that the two changes described above would allow marking to be shown was provided by the work of Fuller (1981) and D.A. Lieberman (personal communication, 1981). They examined spatial discrimination learning in pigeons using the same apparatus as in this research. With a 6-sec delay of reward and a 0.75-sec offset of the keylight alone as a marker, they obtained a higher level of learning in the marked than in the control group.

Method

Subjects

The subjects were 22 adult Rock pigeons which were run in two squads. The first squad of 12 birds had been used in Experiment 4, and the second squad of 10 birds was experimentally naive. All subjects were housed individually in a different room from the experimental and control apparatus, and were maintained at 80% of their pre-feeding weight throughout the experiment. Water and grit were continuously available in the home cages.

Apparatus

The apparatus was the same as that used to run the second squad of subjects in Experiment 4.

Procedure

Magazine training, autoshaping and pretraining. No magazine training, autoshaping or pretraining was necessary for the subjects that had participated in Experiment 4. For the experimentally naive birds the same procedures were followed as in Experiment 4.

Delayed reward training. The subjects were divided into two groups of 11 birds each. The birds in the first squad, (those which had participated in Experiment 4), were run in the same operant chamber as before. If in the last experiment they had been rewarded for choosing green, they now received reinforcement for red choices and vice versa. Allocation of these subjects to the current experimental groups was done in such a way as to match the groups for experimental background, for level of correct responding on the last session of Experiment 4, and for original levels of colour preference.

The naive birds which made up the second squad were allocated to the two groups using the same procedure as in Experiment 4.

The procedure during delayed reward training, which is illustrated in Figure 9, was as follows. Trials began with the illumination of the response key with the discriminative stimuli. In the marked group, the choice response resulted in the immediate offset of the keylight and houselight, both for 0.5 sec. The key was then reilluminated with the red and green stimuli for a further 2.5 sec before changing to white illumination for the final 3 sec of the delay. If the choice response had been correct, the subject then received 3-sec access to wheat grain reinforcement.

In the control group, the choice response simply initiated the trial. There was no change in houselight illumination and the key remained red and green for a further 3 sec. As in the marked group, the key was illuminated entirely white for the final 3 sec of the delay and food reinforcement was then delivered if appropriate.

Birds in both groups received 20 daily sessions of delayed reward training. Each session consisted of 50 trials and the interval between the end of one trial and the next presentation of the red and green discriminative stimuli was 10 sec. Each trial lasted 6 sec from the time of the subject's first, or choice peck to either of these stimuli.

Results

Figure 10 shows the percentage of correct choice responses for each of the first two sessions and averaged over blocks of two sessions thereafter for both groups. There was an improvement in performance for both groups over the course of delayed reward training

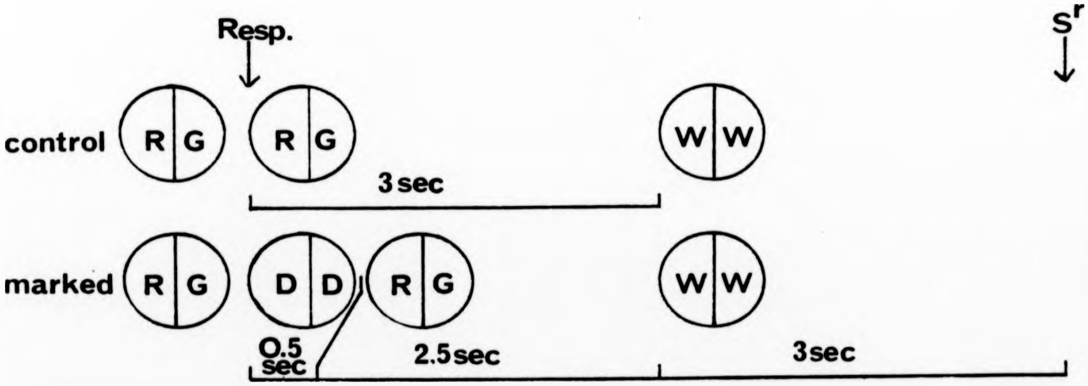

Figure 9

The design of Experiment 5.

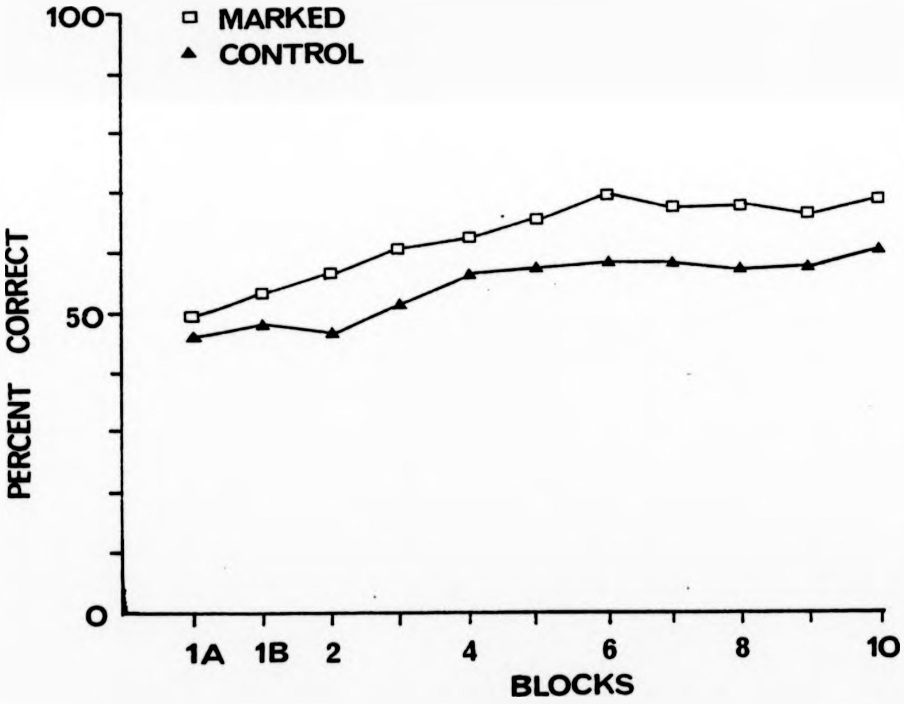

Figure 10

Mean percentage of correct choice responses over the first two sessions and subsequent blocks of two sessions for each group in Experiment 5.

with the marked group showing a greater increase in correct choices and reaching a higher asymptote.

Because the two squads of subjects that were used in this experiment differed in their background - one was naive, the other had been run in Experiment 4 - it was felt necessary to check that their performance did not differ significantly. Accordingly an analysis of variance was performed on the data in blocks of two sessions using three factors - squad, experimental treatment and blocks. It was confirmed that the two squads did not differ, $F(1, 18) = 0.98$, $p > .05$; that there was no interaction between squad and group, $F(1, 18) = 0.09$, $p > .05$; that the Squad × Block interaction was insignificant, $F(9, 162) = 0.23$, $p > .05$; and finally that the Squad × Group × Block interaction was also insignificant, $F(9, 162) = 0.08$, $p > .05$.

The improvement in performance during the experiment that the graph suggests is confirmed by the analysis of variance which revealed a significant change in percentage correct over blocks, $F(9, 162) = 4.32$, $p < .01$. The improvement in correct responding, (measured by subtracting the average figures for the first block from those of the last block), was five percentage points greater in the marked than in the control group and its performance was consistently better. However, analysis of variance showed no significant difference between groups in overall correct responding, $F(1, 18) = 1.41$, $p > .05$, and and no significant Block × Group interaction, $F(9, 162) = 0.21$, $p > .05$.

Discussion

The aim of this experiment was to see whether a reduction in the delay of reward and the use of a new marker would produce a marking effect. In many ways the result was encouraging with the marked group performing consistently, though not significantly, better than the control group. Three interpretations of the outcome of the experiments are available. First, that the effect is real but that the parameters of the design were not optimal. Second, that marking did take place but various other elements of the experimental situation militated against the emergence of a significant difference between the groups. The third interpretation is that the difference between the groups can be attributed to chance. Each of these interpretations will be evaluated in turn.

The first view, that with more suitable parameters an effect would be obtained, has some plausibility. A shortening of the delay of reward and the introduction of a shorter, "flash-type" marker produced a more substantially superior performance by the marked group than in previous experiments. However, it seems unlikely that the failure to obtain a significant effect can be attributed to the level of difficulty of the discrimination. The task was not too difficult - learning was found - nor was it so easy that learning was possible at a rate that would obscure any marking effect.

The second interpretation of this experiment suggests that elements of the design counteracted a marking effect. These elements could have either weakened the performance of the marked group, or promoted learning in the control group, the net effect being to obscure marking. Each element, and the mechanisms through which it may have affected learning, is discussed below.

Adventitious marking

Two potential explanations of the failure to obtain significantly better performance in the marked group are suggested by a consideration of the role played by the change in keylight illumination from red and green to white after 3 sec of the delay. One possibility, (examined in the introduction to Experiment 3), is that this change acted as a non-differential, adventitious marker for both groups. This hypothesis rests on three assumptions - that the birds continued to peck following choice responses; that they did so with an inertial tendency; and that if two responses were marked during the delay interval, then the second one was most likely to be identified as a result of the backward memory scan initiated by reinforcement. If these assumptions are correct, then the effect of the unintended, adventitious marker produced by the change in key illumination would have been to reduce learning in the marked group below the level it would otherwise have reached. This is because the response that was marked adventitiously was, (unless the intertial tendency was complete), sometimes different from the choice response. Thus, marked subjects did not receive consistent marking of the choice response. On the other hand, subjects in the control group benefited from adventitious marking. Whereas they would otherwise have received no marking, they now had responses which were the same as their choice responses occasionally marked. Therefore the net effect of the change in key illumination from red and green to white would have been to narrow the difference between the groups by lowering performance in the marked group and enhancing that of the control group.

Interference hypothesis

The second account of the operation of the change in illumination midway through the delay focuses on the way in which this change may have undermined facilitation produced by the marker in the marked group. This account is suggested by the work of Grant and Roberts (1976) and Wagner (1976, 1978). Grant and Roberts carried out a number of studies of the impact on matching of stimuli interpolated between the sample and comparison stimuli in a delayed matching-to-sample task. They found that such stimuli produced retroactive interference and reduced matching performance. The change from red and green stimuli to white or the addition of the white illumination itself may have acted in a similar way in this experiment. It may have displaced from the birds' short-term memory the visual information about the first response that is required for the solution of the problem (Wagner, 1976, 1978).

However, the difficulty with an account of this kind is that it does not differentiate between the groups, both of which experienced the change in key illumination. Therefore, although this interference interpretation may help explain why performance did not generally improve by a greater amount, it does not provide an explanation of the absence of a significant difference between groups.

One factor that might have undermined a marking effect in this experiment was the use of houselight offset as a marker element. The effect of this event might have been to draw the subject's attention away from the response key and cause it to look around. One way in which distraction of this kind might have damaged marking is suggested by the interference hypothesis discussed above. Attention to other stimuli in the chamber could have disrupted processing of the choice response and thereby reduced the benefits of marking. Some evidence against this possible interpretation is provided by Lieberman et al.'s (1979) successful use of white-noise bursts and light flashes as markers, events which might also be seen as potentially interfering distractors.

Distractor hypothesis

There is also a second way in which a distractor effect of houselight offset might have undermined marking in this experiment. One feature of learning in this situation may be the continued availability of the discriminative stimuli after the choice responses. This might have helped in two ways, first, through the consolidation of stimulus processing, and second, through repeated responses to the chosen stimulus as a result of an intertial tendency in pecking. If houselight offset distracted marked subjects' attention away from the stimuli, then it would reduce the benefits of continued exposure to the stimuli. Such benefits would, however, be available to control subjects. Evidence to support the view that learning may be facilitated by the continued availability of the stimuli is provided by Hudson in an experiment reported by Tolman (1948, p.200). Hudson examined one-trial avoidance learning in rats in a situation where electric shock was associated with a visual pattern. A strong avoidance response to the pattern was found except when the pattern was removed immediately following the shock, suggesting that continued exposure to it assisted learning.

Evidence against the view that the continued availability of the stimuli enhances learning is provided by two experiments. The first is Thomas et al.'s (1983) final experiment where they ensured that no discriminative cues were available at the time of marker presentation after a choice response. Although their experiment did not involve a comparison of performance with and without the continued presence of cues, it could be argued that their result implies that it should not matter if the subject's attention is taken away from the stimuli, (or if the stimuli are removed immediately after a choice response). Nonetheless it could still be argued that there are potentially important differences between the maze and operant chamber situations which allow a distractor effect in the latter case but not the former. For instance, the subjects in maze experiments were not able to repeat their initial choice response in the way that subjects in the operant chamber could. However, the proposition that a distractor effect benefited the control group may be difficult to sustain in the face of evidence from Experiment 4. Two of the groups from that experiment are of particular interest - the control and whole-key-marker groups. In the control group, choice response simply initiated the trial and the red-green stimuli remained on the key for a further 3 sec, whereas choices by marked subjects darkened both sides of the key for the same period. Learning was the same in both groups which suggests that continued exposure to the discriminative stimuli does not help control subjects and therefore that a distractor effect, if it exists, should not harm marked subjects. This conclusion is further strengthened by direct evidence from the current experiment. If there was a distractor effect of sufficient strength to damage marked group learning then one might expect to find a lower response rate in that group during the first 3 sec of the delay (before the key illumination changed to white). Although the marked group's response rate averaged over training sessions was indeed lower - 372 responses per session as opposed to 450 for the control group - the difference was not significant, $t(20) = 0.57$, $p > .50$. Even on the first session when any distractor effect should have been at its strongest given the novelty of houselight offset at that stage, the difference between the rates for the marked and control groups, 257 and 431 responses per session respectively, was not significant, $t(20) = 1.37$, $p > .20$. Evidence, therefore, that the continued availability of the discriminative cues after a choice response might facilitate learning or that marked subjects failed to receive this benefit because they were distracted by the houselight offset is weak.

Summary and conclusions

So far two interpretations of the outcome of this experiment have been reviewed. Both of them assume that the difference between marked and control groups reflected a real effect. One interpretation was that the selection of parameters was not ideal, the other that the presence of certain elements in the design militated against a marking effect. The argument that with a more optimal choice of parameters a significant effect would emerge has some plausibility. However, the identification in this discussion of design elements that could have conferred disadvantages on the marked group and advantages on the control group suggests that more weight should be given to the second interpretation. One possibility which has been discussed is that the change in key illumination during the delay interfered with performance in the way suggested by Grant and Roberts (1976) and Wagner (1976, 1978) for other experimental situations. This interference interpretation can explain why learning was not better in both groups but does not differentiate between them. However, a second account of the role of key-illumination change is more useful. If the change acted as an adventitious marker, then learning in the marked group would be reduced below the level it would otherwise reach and control group learning would be enhanced. Hence, a narrowing of the difference between the groups would occur. A third way in which the extent of a marking effect could be undermined was suggested by a consideration of the impact of the houselight-offset element of the marker. This might have caused marked subjects to look around the chamber and attend to other stimuli, thereby interfering with processing of the choice response and damaging learning. A final possibility was also suggested by the offset of houselight. If the continued availability of the discriminative stimuli following choice responses facilitated learning, then another consequence of a distracting houselight offset would be to undermine learning in the marked group. However, this possibility received limited support.

The third interpretation of the result of Experiment 5 was that the superior performance of the marked group did not reflect a real, marking effect but could simply be ascribed to chance. This is clearly a possibility, but given the chance that the parameters of the experiment were not ideal and, more importantly, that several factors have been identified which might have militated against a marking effect being revealed, it seems premature to conclude at this stage that the effect was not real.

3.2 EXPERIMENT 6

The outcome of the last experiment was the most encouraging so far in this series of experiments with the marked group performing consistently, though not significantly, better than the control group.

The purpose of this experiment was to continue to look for marking using a modified procedure, (which is shown in Figure 11), that took into account some of the potential difficulties raised by Experiment 5. The change in that experiment to a brief "flash" of a marker appeared to have been beneficial and it was therefore decided to use a very brief, 0.5 sec, marker in this experiment as well.

One factor that should be important in determining the effectiveness of a marker is its novelty. With the present design birds clearly had substantial experience of a darkened response key both during pretraining and in the intertrial intervals of the training stage. The darkening of the keylight immediately following a choice response in these experiments was therefore less novel than would be a change to an entirely new colour. Accordingly in this experiment, the key was illuminated with amber following a choice response rather than darkened.

A second change in design for this experiment was the discontinuation of houselight offset as an element of the marker. As the discussion of the last experiment indicated, this event might possibly function as a distractor which reduced the effectiveness of the marker and hindered rather than helped attempts at marking.

In order to test the hypothesis that the change in key illumination from red and green to white might have undermined a marking

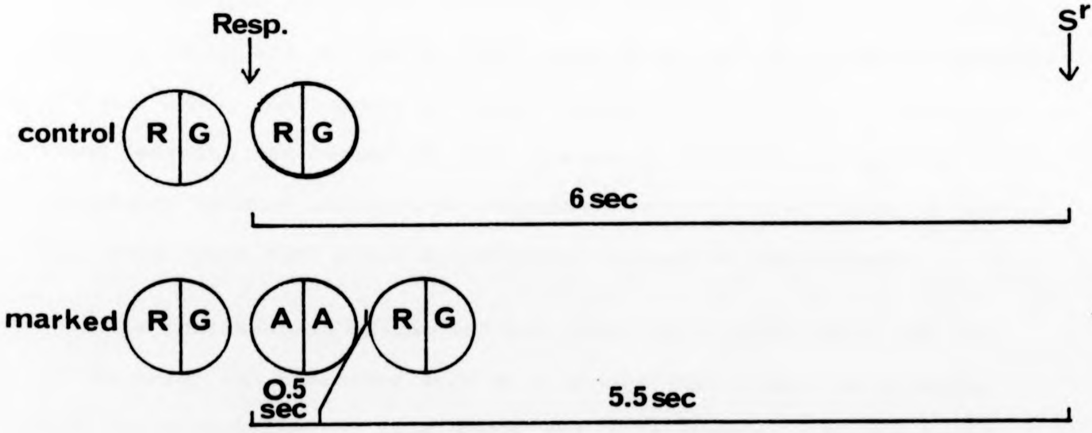

Figure 11

The design of Experiment 6. (A = amber.)

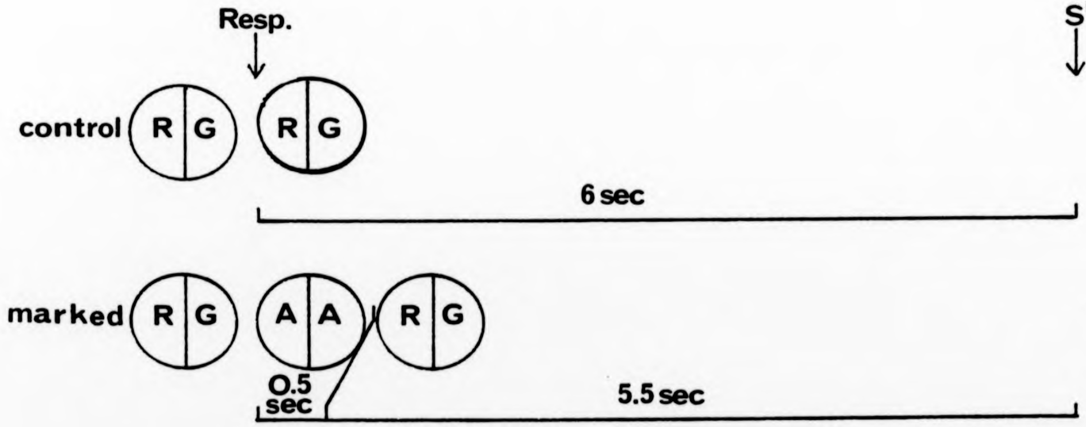

Figure 11

The design of Experiment 6. (A = amber.)

effect, another revision was made to the general design. In the current experiment the red and green lights remained until the end of the delay. The advantages of this amendment were several - one, it eliminated a possible adventitious marker; two, it reduced any interference with visual processing and/or memory because of the additional processing required by the white illumination; and three, it provided an aid to learning by maintaining the same cues at the beginning and end of the delay. (The third advantage is suggested by Spear's (1973, 1978) reinstatement theory which was discussed in the first chapter of the thesis.) A disadvantage anticipated with this change was that it increased the chances of adventitious learning. If subjects continued to peck the response key and did so with an inertial tendency, then chance pairings of responses to the correct colour with food could occur. A particular concern was that there would be an opportunity for the control group to improve at a rate that would obscure any marking effect. However, since the beginning of this research programme, experiments conducted involving a spatial discrimination had not encountered major problems with control subjects (Fuller, 1981; Lieberman, personal communication, 1981). Adventitious learning had seemed an even greater possibility in spatial situations because pairings of responses to the correct side of key with food could take place despite the removal of the stimuli.

The continuation of the red and green key-lights until the end of the delay was therefore seen as a change that would, on balance, make the discrimination task easier for both groups. It was hoped that it would provide a more sensitive design which was less susceptible to a low ceiling effect and more likely to reveal marking.

If the marking phenomenon is a general one then the use of an entirely novel change to an amber keylight immediately following choice responses should produce better learning in the marked than in the control group.

Method

Subjects

The subjects were 12 adult Rock pigeons which were experimentally naive at the start of the experiment. They were housed and maintained under the same conditions as in earlier experiments.

Apparatus

The apparatus was the same as in Experiment 5.

Procedure

Magazine training, autoshaping and pretraining. The same procedure was used as in Experiment 4.

Delayed reward training. The birds were divided into two groups of six subjects each. The groups were matched for level of colour preference, level of side preference, and response rate on the basis of the birds' performance in the pretraining session. The running order for subjects and their allocation to each of the two operant chambers were balanced across groups. The colour to be rewarded was determined in the same way as in Experiment 4.

A trial began with the illumination of the response key with the discriminative stimuli. The subject's first peck represented its choice response. Subjects in the marked group were presented with a change in keylight illumination from red and green to amber for 0.5 sec immediately following a choice response. No other markers occurred. After 0.5 sec of amber, the key was reilluminated with red and green lights which remained available for a further 5.5 sec until the end of the 6-sec delay. Subjects in the control group simply initiated the delay interval with their choice response and the red-green illumination stayed on for the entire 6 sec of the delay.

Subjects in both groups received 3-sec access to wheat grain at the end of the trial if their choice response had been correct. The interval between the end of one trial and the next presentation of the discriminative stimuli was 10 sec. During this interval the key was darkened. All subjects were given 20 daily sessions of 50 trials.

Results

The percentage of correct choice responses averaged over blocks of two sessions for both groups is shown in Figure 12. The figure shows that although both groups improved over the course of the epxeriment, this improvement was greatest for the control group which had reached 80% correct by the final block, compared with 60% for the marked group.

A two-way analysis of variance for experimental treatment and block was performed and showed that the control group made significantly more correct responses than the marked, $F(1, 10) = 5.79$, $p < .05$. There was also a significant increase in percentage correct during delayed reward training, $F(9, 90) = 10.98$, $p < .01$. Although the Blocks × Group interaction was not significant, $F(9, 90) = 0.54$, $p > .05$, the control group appeared to improve from a similar starting point to a substantially higher asymptote than the marked group.

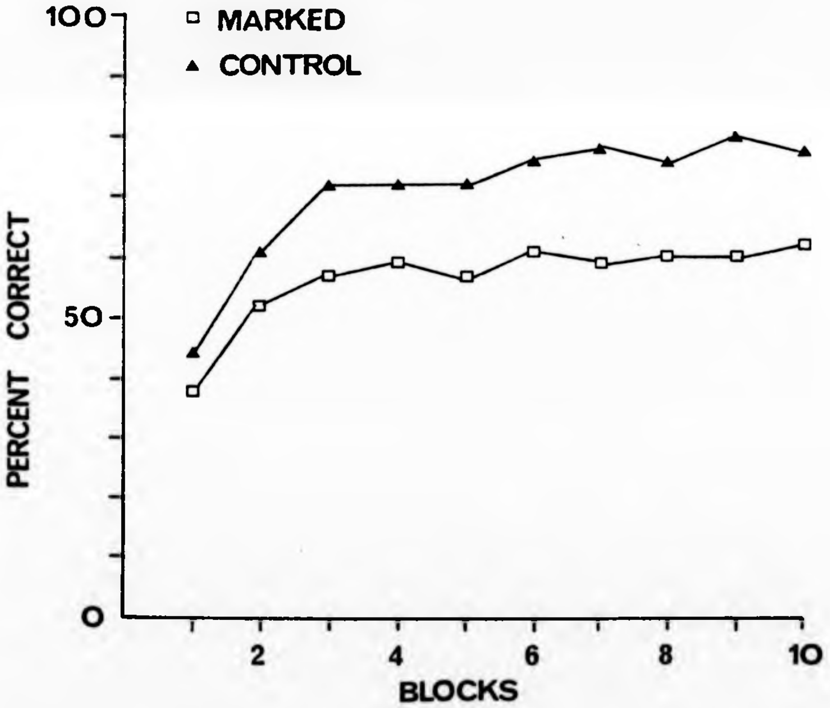

Figure 12

Mean percentage of correct choice responses over blocks of two sessions for each group in Experiment 6.

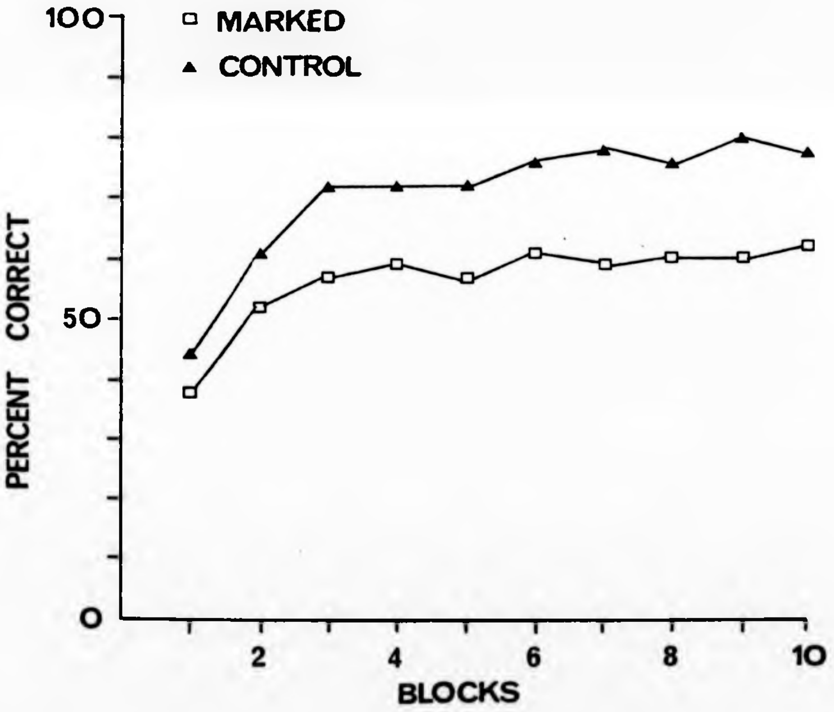

Figure 12

Mean percentage of correct choice responses over blocks of two sessions for each group in Experiment 6.

Discussion

The intention in this experiment was to modify the relatively successful design of Experiment 5 in ways which might provide appropriate conditions for revealing a marking effect. No such effect was found - rather the level of correct responding was significantly greater in the control group than in the marked group. A result of this kind, which not only fails to support the marking hypothesis but actually runs counter to it, clearly raises several questions.

One issue is why the control group performed so well both in relation to the marked subjects and in comparison with the control groups of Experiments 4 and 5. In Experiment 4 the percentage of correct responses by the control group rose from 42% to 58% over the course of delayed reward training, and the comparable figures for Experiment 5 were 47% and 61%. In the current experiment an average increase for control subjects from 44% to 78% was observed. A new feature of the design used this time was that the red and green discriminative stimuli remained on the response key until the end of the delay rather than being replaced with white stimuli after 3 sec. One possibility, which was raised in the discussion of Experiment 5, was that this stimulus change interfered with memory for or processing of the discriminative cues or choice response. Its removal in the current experiment may therefore have helped learning.

An additional possibility is that the maintenance of the discriminative stimuli until the end of the delay allowed adventitious reinforcement to occur in a way that was not possible in the earlier experiments where the stimuli were no longer present at the time of reinforcement. If the subjects continued to respond and did so with an inertial tendency, then learning through adventitious reinforcement could occur. Observation of the birds' behaviour and data collected provide clear evidence that at least the first condition is met.

A more difficult issue raised by this experiment concerns the performance of the marked group. Not only did this group do no better than the control group, its level of learning was actually significantly worse. The poorer performance by the marked group is especially surprising for four reasons. First, marked subjects, like those in the control group, should have benefited from the removal of the potentially interfering change in illumination during the delay. Second, the marked group as well as the control group should have gained from adventitious learning made possible by the maintenance of the discriminative stimuli. Third, there was no off-the-response-key event, such as houselight offset, which could possibly function as a distractor of the subject's attention. Finally, the marker was an entirely novel stimulus. A further problem is that this result conflicts with the findings of the last experiment where the marked subjects did better than the controls although the difference was not statistically significant.

Why then was the level of learning in the marked group lower? One account of the role of stimulus change on the response key has already been proposed in the discussion of the last experiment and in the consideration above of the control group's performance. This interference account suggests that the impact of either stimulus change or number of stimuli might have been to displace information from the birds' short-term memories.[1] Consequently the birds would

---

[1] The difference between stimulus change and number of stimuli in this context can be illustrated by considering the keylight sequences red-green-blue and red-green-red. Both involve two changes but the first sequence involves one extra colour. The relative importance of these factors may be different and could, of course, be tested experimentally by comparing the consequences of various delay interval sequences.

not have had the information necessary for the solution of the problem available at the end of the delay. An account of this kind is suggested by Wagner's model of the impact of surprising events on rehearsal (Wagner, 1976, 1978; Wagner et al., 1973). In the discussion of Experiment 5 this account was invoked in an exploration of the roles of the change in key illumination and of houselight offset as a marker element. Given the results of the current experiment it may be useful to extend the analysis to include the marker itself as a potentially interfering event.

The merits of the interference account could be evaluated in two ways. The first is by a comparison of the performance of the marked groups in Experiments 5 and 6. In the last experiment the marked group received both more stimulus changes and more stimuli than the marked group in this experiment. The former subjects experienced the response key with red and green, white, and no illumination as well as houselight offset, whereas the later subjects were exposed to amber alone, and to red and green illumination. The group which was exposed to most stimulus change - that of Experiment 5 - improved slightly less and at a slower rate although its asymptotic level was higher. However, interpretation of this slight difference is made even more difficult by the fact that the marked group in the current experiment had the opportunity to benefit from adventitious reinforcement. Thus a comparison of the marked groups does not provide any clear evidence for or against an interference view.

A second way of evaluating this account is to consider the relative positions of the marked and control groups in each of these experiments. In Experiment 5 the marked group experienced more potentially interfering events than the control group. In addition to stimulus changes received by both groups, such as the switch from red and green to white key illumination, the marked group also experienced the houselight offset and keylight offset which made up the marker. An interference interpretation would therefore predict that performance should be worst in the marked group. However, in Experiment 5 the marked group in fact did better than the control group. In Experiment 6 the opposite result was found - the marked group, (despite less stimulus change than its counterpart in the last experiment), showed a significantly lower level of correct responding than the control group. Thus the outcome of Experiment 5 on its own, and the results of the last two experiments when considered together do not provide support for a coherent interference interpretation of the impact of stimulus change or number of stimuli on visual memory.

One way of reconciling the different outcomes of this and the last experiment within an interference hypothesis might be to focus on the characteristics of the markers used in each case. In Experiment 5, the marker was a complete darkening of the operant chamber with both keylights and houselight being switched off. An event of this kind may be less damaging to the processing of information about colour than the amber marker used in this experiment. A more specific version of the interference account would suggest that the level of retroactive interference is dependent upon whether or not the discriminative stimuli and the marker are drawn from the same class. A direct comparison of these two markers within a single experiment would be helpful in giving more conclusive evidence about their relative effects on learning.

Another dimension to the question of why the marked group fared so badly in this experiment relates to its apparent failure to benefit from adventitious learning. The continued presence of the discriminative stimuli did appear to help the control group, possibly through an inertial response tendency allowing adventitious reinforcement. The operation of the marker offers two potential explanations of why subjects in that group did not do at least as well as the controls. If the interference hypothesis is correct, the disruption of colour-processing may persist until the end of the delay and prevent learning about an adventitiously rewarded response to the correct stimulus. On the other hand, the presentation of the marker may have produced a "startle effect" which disrupted any inertial tendency, thereby removing the basis for adventitious learning.

In summary, the substantial learning produced in the control group of this experiment may have been a result of adventitious learning. The cause of the poor performance of the marked group is less clear - one possibility is interference with colour memory, another is disruption of inertial responding by the marker.

3.3 EXPERIMENT 7

The aim of this experiment was to address some of the issues raised by the last two experiments. One of these was the role that an inertial tendency in the pecking behaviour of the subjects might have played in learning. An examination of the birds' response patterns would be helpful in two ways. First, if an inertial tendency was found in the control group it would provide support for the view that the improved performance by that group in the last experiment might be attributable to adventitious learning. Second, if such a tendency was found in the control group but not in the marked group, then this would constitute evidence that one reason for the poor learning in the marked group was that the marker disrupted the basis for adventitious reinforcement. One purpose of this experiment, therefore, was to collect data which would provide information about any inertial tendencies in the subjects' response patterns.

Another major consideration in the design of this experiment was the desirability of replicating Experiment 6. This was seen as particularly important on two counts - one was that its outcome had been directly counter to that anticipated on the basis of the marking hypothesis. The other was that the results conflicted with the findings of Experiment 5 in which the marked group had done better, (although not significantly so), than the control group. An additional benefit of the design was that it allowed a direct comparison of the consequences of different markers within a single experiment.

Two marked groups were used. In one, described as the blackout-marker group, choice responses were followed by the offset of both the houselight and keylight for 0.5 sec. Subjects in the other group

- called the amber-marker group - had their choices followed by the illumination of the key with amber light for 0.5 sec. In both groups the response key was red and green for the final 5.5 sec of the 6-sec delay. Thus the markers which had been used in Experiments 5 and 6 were brought together in a single design. The only modification to the groups used previously was in the case of the blackout-marker group where in Experiment 5 the keylights had changed from red-green to white after 3 sec. A control group that experienced no markers was also run.

Method

Subjects

The subjects were 15 adult Rock pigeons which were experimentally naive at the start of the experiment. They were housed and maintained under the same conditions as in earlier experiments.

Apparatus

The apparatus was the same as in Experiment 5.

Procedure

Magazine training, autoshaping and pretraining. The same procedure was used as in Experiment 4. Some subjects required two sessions of autoshaping or pretraining.

Delayed reward training. Three groups of five birds each were formed. Each group was matched with the others for level of colour preference, level of side preference, and response rate on the basis of the birds' performance in their last session of pretraining. The allocation of subjects to each of the operant chambers and their daily running order were balanced across groups. The colour to be rewarded was determined in the same way as in Experiment 4.

A trial began with the illumination of the response key with the discriminative stimuli. The subject's first peck after the start of the trial was designated its choice response. For birds in the blackout-marker group, choice responses were immediately followed by the offset for 0.5 sec of both the houselight and the keylight. Subjects in the amber-marker group had choices immediately followed by amber illumination of the response key for 0.5 sec. In both these groups, after the marker the key was reilluminated with red and green lights which remained on for the final 5.5 sec of the 6-sec delay. Subjects in the control group simply started the trial with their choice response and the red and green stimuli were present for the entire delay interval.

Correct choices were rewarded in all groups by 3-sec access to wheat grain at the end of the trial. The interval between the end of one trial and the next presentation of the discriminative stimuli was 10 sec. During this interval the response key was dark. All subjects received 20 daily sessions, each containing 50 trials.

Results

The primary data shown in Figure 13 are the percentage of correct choice responses averaged over blocks of two sessions for each of the groups. It can be seen that the control group showed considerable improvement, increasing from approximately 40% correct at the start to around 75% by the end of training. The performance of the blackout-marker group was similar though it rose to an asymptote about were balanced across groups. The colour to be rewarded was determined in the same way as in Experiment 4.

A trial began with the illumination of the response key with the discriminative stimuli. The subject's first peck after the start of the trial was designated its choice response. For birds in the blackout-marker group, choice responses were immediately followed by the offset for 0.5 sec of both the houselight and the keylight. Subjects in the amber-marker group had choices immediately followed by amber illumination of the response key for 0.5 sec. In both these groups, after the marker the key was reilluminated with red and green lights which remained on for the final 5.5 sec of the 6-sec delay. Subjects in the control group simply started the trial with their choice response and the red and green stimuli were present for the entire delay interval.

Correct choices were rewarded in all groups by 3-sec access to wheat grain at the end of the trial. The interval between the end of one trial and the next presentation of the discriminative stimuli was 10 sec. During this interval the response key was dark. All subjects received 20 daily sessions, each containing 50 trials.

Results

The primary data shown in Figure 13 are the percentage of correct choice responses averaged over blocks of two sessions for each of the groups. It can be seen that the control group showed considerable improvement, increasing from approximately 40% correct at the start to around 75% by the end of training. The performance of the blackout-marker group was similar though it rose to an asymptote about

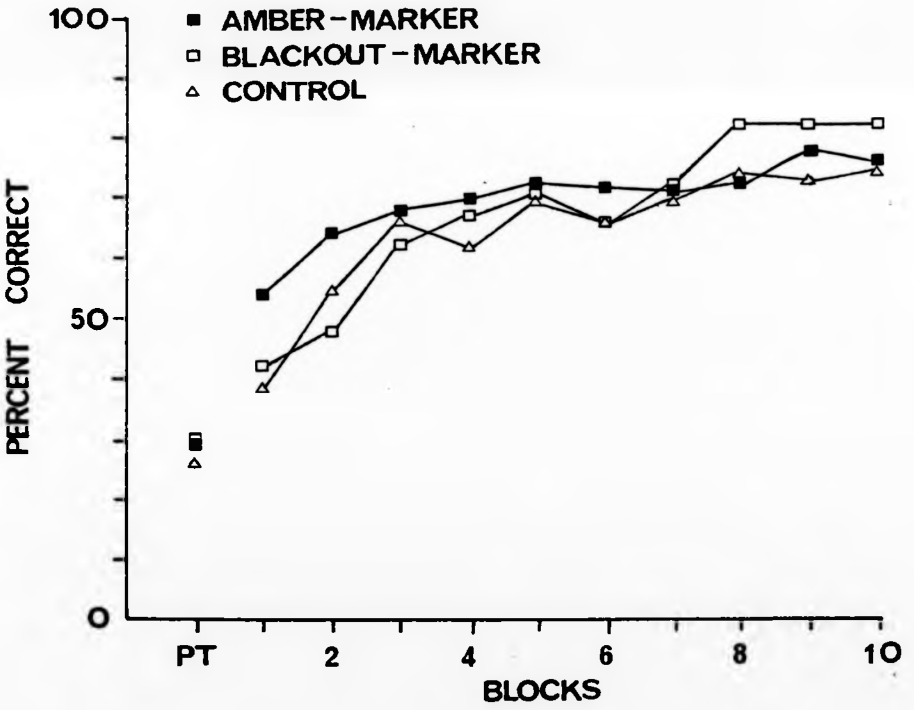

Figure 13

Mean percentage of responses to "correct" colour during the final session of pretraining and percentage of correct choice responses over blocks of two sessions for each group in Experiment 7.

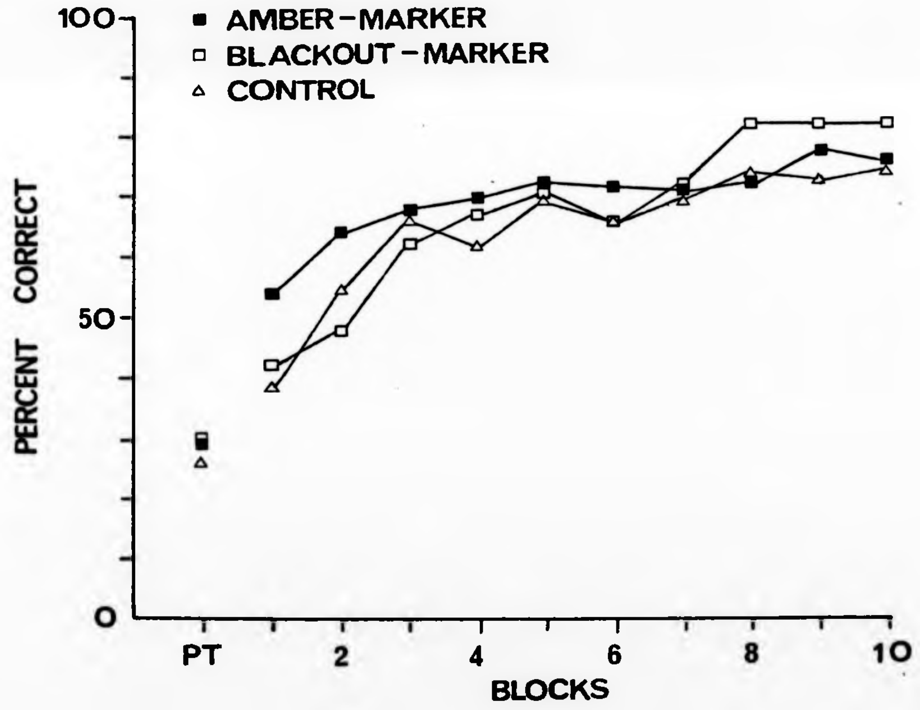

Figure 13

Mean percentage of responses to "correct" colour during the final session of pretraining and percentage of correct choice responses over blocks of two sessions for each group in Experiment 7.

10 percentage points above that of the control group. The results for the amber-marker group are difficult to compare with those of the other two groups as its starting point was higher. Its level of improvement was less than that of either the blackout-marker or control group, with a rise of about only 20 percentage points.

A two-way analysis of variance using the factorsof experimental treatment and blocks was performed and this revealed that there was indeed a significant increase in percentage correct over the course of the experiment, $F(9, 108) = 8.2$, $p < .01$. No difference in levels of correct responding between groups was found, $F(2, 12) = 0.19$, $p > .05$, and the Blocks × Group interaction was also insignificant, $F(18, 108) = 0.50$, $p > .05$.

The source of the difference in starting points for the groups, (which is apparent on the first session of delayed reward training), is unclear. Figure 13 also shows the percentage of correct responses to the subsequently - designated, correct colour on the last session of pretraining for each group. It is clear from the figure that the groups did not differ substantially in their level of colour preference prior to the start of delayed reward training and this is confirmed by analysis of variance, $F(2, 12) = 0.11$, $p > .05$.

Discussion

This analysis will look first at the performance of the control and blackout-marker groups in this experiment. The control group again showed a very substantial increase in percentage correct, replicating very closely in both pattern and level of learning the results found in the same group in Experiment 6. This result, which is illustrated in Figure 14, shows that if the discriminative stimuli remain available until the end of the delay, then for control subjects at least, learning is enhanced compared with situations where the stimuli are removed earlier. A strong possibility is that an inertial tendency in responding could have led to adventitious reinforcement and learning. Unfortunately an error in data collection meant that the evidence about pecking behaviour which might support this interpretation did not become available.

The blackout-marker group also did well with a considerable improvement over the course of training. Maintenance of the discriminative stimuli during the delay interval appeared to help the blackout-marker group in a similar way to the control group. In Experiment 5 - where the red and green stimuli changed to white after 3 sec - the marked group rose from about 50% correct on the first training block to around 70% on the final block. In this experiment the figures were around 40% and 80% respectively. Thus both the blackout-marker groups and control groups discussed above benefited by approximately 25 percentage points from the continued presence of the discriminative stimuli.

Taking the control and blackout-marker groups together suggests a pattern that is at least consistent with earlier results and is not directly counter to the marking hypothesis. Figure 13 shows that in the final stages of the experiment, the performance of the blackout-marker group was marginally superior to that of the control group albeit not significantly so. The difference between the groups matches that found in Experiment 5 with groups of similar design.

The interpretation of the amber-marker group is more difficult. The pattern of its performance in terms of the level and rate of

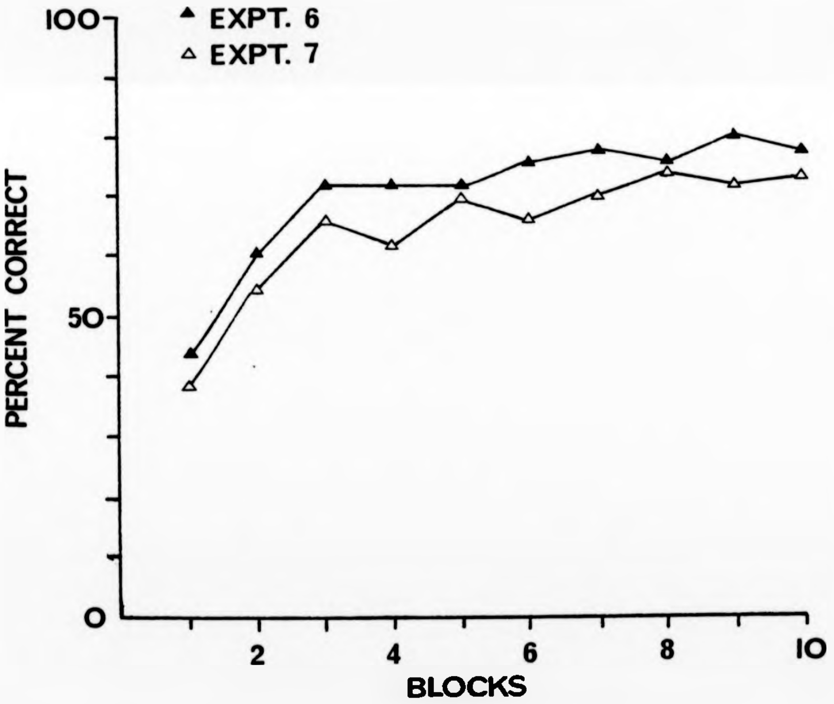

Figure 14

Mean percentage of correct choice responses over blocks of two sessions for the control groups in Experiments 6 and 7.

learning is very similar to that of the identical group in Experiment 6. This can be seen in Figure 15 where the learning curves of amber-marker groups from Experiments 6 and 7 have been plotted with a shared starting point. In both experiments the improvement in percentage points was about 20. However, in the present experiment the starting point for the amber-marker group was quite different from that of both the blackout-marker and control groups, and the comparable group in the last experiment. This difference in starting points is difficult to account for. As can be seen from Figure 13, the groups in Experiment 7 were well matched on the basis of their pretraining performance. The difference may therefore be a result of either random variation or an immediate marking effect producing very rapid early learning. The latter seems unlikely primarily because no such effect was found in the last experiment where, on the first block of training, the performance of the marked group was marginally below the control group's. It is also the case that the asymptotic learning level reached by the amber-marker group is both slightly less than that of the blackout-marker group and is reached marginally later.

Two approaches can be adopted to the results of the amber-marker group. The first is to replot the learning curve so that its starting point is close to those of the other two groups. If this is done then the outcome of this experiment is very like that of the last one with the control group doing substantially better than the amber-marker group. Taken in conjunction with the blackout-marker group, this suggests that there is something about the nature of the amber marker which is particularly damaging to colour discrimination learning. The version of the interference hypothesis which points to the

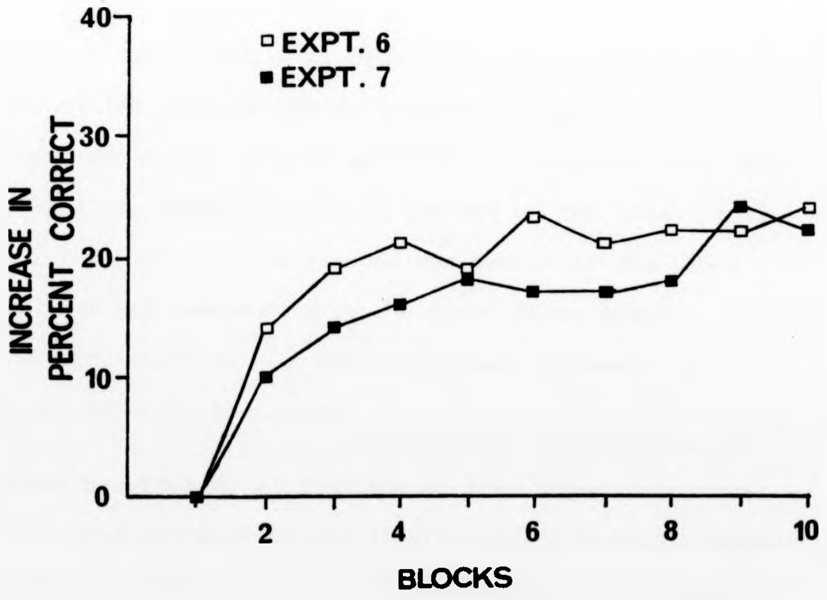

Figure 15

Increase in percent correct choices over blocks of two sessions for the amber-marker groups in Experiments 6 and 7.

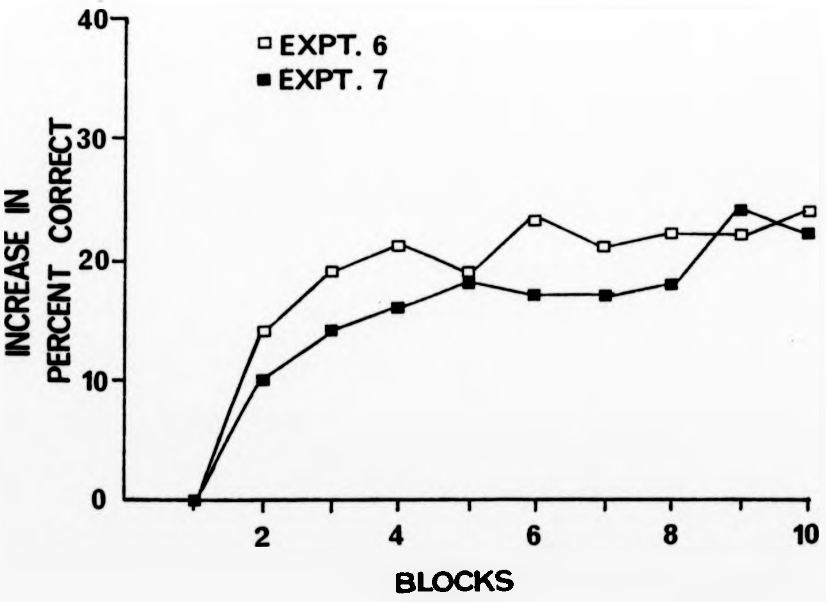

Figure 15

Increase in percent correct choices over blocks of two sessions for the amber-marker groups in Experiments 6 and 7.

greater capacity of the amber marker than the blackout marker to interfere with colour memory is therefore supported.

On the other hand, if the learning curve shown by the amber-marker group is not replotted, then the effect of this marker is unclear. The results of Experiment 6 are not replicated and one is left with a situation where the two experiments point to different conclusions. The first showed that the amber marker was damaging and reduced performance below that of the control group, whereas the second implied that the marker probably had no impact as the asymptotes in the marked and control groups were the same. What is clear from both experiments, (unless the initially high level of the group in this experiment is ascribed to marking), is that there was no enhancement of learning when choice responses were followed by an amber marker.

The analysis of Experiment 7 is made difficult on two counts. One, the absence of the appropriate data makes it impossible to be entirely confident about the role of adventitious learning when discriminative stimuli are maintained until the end of the delay. It is not possible, for instance, to examine the impact of the amber marker upon any inertial tendency that the birds might display. The second source of difficulty is the ambiguity that surrounds the performance of the amber-marker group.

Despite these problems it is possible to draw three tentative conclusions - one, the presence of the discriminative stimuli throughout the delay allows a higher level of learning; two, the blackout-marker produces results which are at least as good as the control groups and do not strongly contradict the marking hypothesis; and three, that a specific, colour interference hypothesis may be consistent with these results.

3.4 EXPERIMENT 8

The basis for the last three experiments was an attempt to demonstrate a marking effect and thereby establish that earlier failures to do so reflected inappropriate choices of parameters and not some more fundamental limitation on the generality of the phenomenon. However, no significant marking effect was found in Experiments 5 to 7 although the results did provide some useful pointers. The aim in this experiment was to incorporate those features in a design which, it was hoped, would reveal a marking effect. A successful demonstration of marking would still leave a number of questions to be dealt with about the reasons for past failures. Nonetheless, evidence that marking was possible using a quite different paradigm to that of the original studies by Lieberman et al. (1979) and Thomas et al. (1983) would be valuable.

One design feature that was identified as being important if marking was to be revealed, was a control group which did not learn too rapidly. Two different control groups had so far been used with a 6-sec delay of reward. The first of these, in Experiment 5, had shown a relatively small increase in percentage correct over the course of training. The group's average percentage correct had only risen from about 50% on the first training block to just over 60% on the final block. In contrast, the second type of control group, (found in Experiments 6 and 7), had performed very well, showing rapid learning and a high asymptotic performance. In both experiments, the control subjects had improved by an average of 25 percentage points during the first two blocks alone. Performance of this kind seemed likely to make it difficult to avoid ceiling effects which would obscure any benefit of the marker. Therefore, it was decided to adopt in this experiment the same type of design as had been used in Experiment 5. Thus in the control group, choice responses started the trial and after 3 sec the discriminative stimuli were replaced with white illumination which remained until the end of the delay.

Of the various marked groups that had been used so far, the most promising results had been obtained in Experiment 5 in which choice responses were followed by the offset of the houselight and keylight for 0.5 sec. The level of correct responding in the marked group was consistently, though not significantly, higher than in the control group. When the same marker had been used in the last experiment, (with a different delay-interval procedure), it also produced very slightly better performance than in control subjects towards the end of training. In the present experiment a new element was added in an attempt to strengthen the marker and make it more likely to initiate a backward scan of memory. The new element was the onset of the magazine light for 0.5 sec immediately following a choice and concurrent with offset of the houselight and keylight. The magazine light should have acquired secondary reinforcing properties through repeated pairings with food during magazine training, autoshaping and pretaining. Consequently its ability to attract the attention of the subject should be high and it should in combination with the other elements, provide a salient and unexpected event of the kind that would be expected to act as an effective marker. Although there was a possibility that magazine-light onset would act as a distractor in the way discussed in Experiment 5, the evidence to support that analysis seemed weak and was considerably outweighed by the potential benefit of adding this element. Following the marker, the red and green stimuli returned to the key and remained on for a further 2.5 sec before being replaced with white illumination until the end of the delay.

Method

Subjects

The subjects were 14 adult Rock pigeons which were experimentally naive at the start of the experiment. They were housed and maintained under the same conditions as in previous experiments.

Apparatus

The same apparatus was used as in Experiment 5.

Procedure

Magazine training, autoshaping and pretraining. The same procedure was used as in Experiment 4. Some subjects required two sessions of autoshaping.

Delayed reward training. The subjects were divided into two groups of seven birds each. The groups were matched for level of colour preference, level of side preference, and response rate on the basis of the birds' performance in pretraining. The allocation of subjects to each of the operant chambers and the daily running order of subjects were balanced across groups. The colour to be rewarded was determined in the same way as in Experiment 4.

A trial began with the illumination of the response key with the discriminative stimuli. The subject's first peck represented its choice response. Subjects in the marked group had choice responses immediately followed by the offset of houselight and keylight, also for 0.5 sec. The response key was then reilluminated with the red and green discriminative stimuli which remained for a further 2.5 sec before being replaced with white illumination. The key stayed white for the final 3 sec of the 6-sec delay.

Subjects in the control group simply initiated the trial with their choice response. The red and green stimuli remained available for 3 sec before replacement with white illumination until the end of the delay.

In both groups correct choices were rewarded by 3-sec access to wheat grain at the end of the trial. The interval between the end of one trial and the next presentation of the discriminative stimuli was 10 sec. During this interval the key was darkened. All subjects received 20 daily sessions, each of 50 trials.

Results

The percentages of correct choice responses averaged over blocks of two sessions for both groups are presented in Figure 16. It is apparent from Figure 16 that both groups show an improvement in correct responding during the experiment and that there is no difference between the average performance of marked and control subjects. This impression is supported by a two-way analysis of variance using the factors of experimental treatment and blocks. There is a significant change in the percentage correct over blocks, $F(9, 108) = 4.14$, $p < .01$, but no significant groups effect, $F(1, 12) = 0.17$, $p > .05$, or Blocks × Group interaction, $F(9, 108) = 0.32$, $p > .05$.

Discussion

The results of this experiment do not support the hypothesis that following a choice response with a salient event will facilitate learning. Despite the use of an apparently optimal combination of

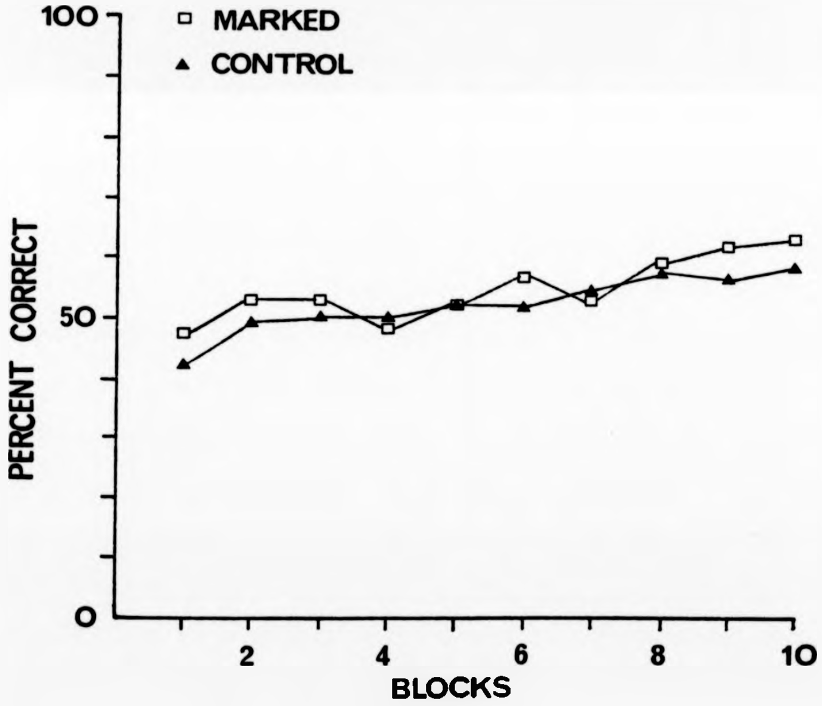

Figure 16

Mean percentage of correct choice responses over blocks of two sessions for both groups in Experiment 8.

elements as a marker, the performance of the marked group was no better than that of the control group.

A comprehensive analysis of the interpretations that can be placed on the results of this and earlier experiments will be provided in the general summary and discussion which follows this section. The discussion here will focus on why the new element of the marker - the onset of the magazine light for 0.5 sec - did not provide the additional benefit that had been anticipated. In Experiment 5, where offset of the houselight and keylight was used as a marker, the marked group performed better, though not significantly so, than the control group. Not only did the extra element fail to enhance this slight effect, but the learning curves of the two groups were actually indistinguishable in the current experiment.

In retrospect the use of a flash of magazine light might have been damaging in three ways. The first is that the effectiveness of the offset of houselight and keylight as a marker in Experiment 5 might have depended on the complete darkening of the chamber which it involved. The partial illumination of the chamber by the magazine light might have served to undermine the potency of the other two elements by reducing the amount of change in chamber illumination following choice responses. A second way in which the added element could have impaired learning in the marked group relative to the control group was by acting as a distractor. Given the birds' experience of magazine light pairings with food, the onset of this light would be very likely to draw the birds' attention away from the response key and towards the magazine. It is possible that control subjects benefited from continued exposure to the discriminative stimuli, (though the discussion in Experiment 5 indicated this was unlikely), and consequently that the addition of magazine light onset to the marker detracted from learning in marked subjects. A third possibility is based on the likelihood that the magazine light had been established as a secondary or conditioned reinforcer by virtue of its repeated pairings with food during the stages before delayed reward training and the training itself (Kelleher, 1961). As a consequence its use as a marker might have had the effect of strengthening both correct and incorrect responses and thereby detracting from any advantage which might have been derived from a marking effect acting on its own.

Unfortunately it is difficult to provide evidence from the present experiment that could test the validity of these three hypotheses. A comparison of response rates in the two groups might be informative about either the distractor or the secondary reinforcer mechanisms if either was operating in isolation. One would expect the magazine light onset to lower the response level in the marked group relative to the control group if it acted as a distractor. On the other hand, if it possessed secondary reinforcing qualities then an enhancement of responding in the marked group would be anticipated. However, as both effects could well have been at work simultaneously no clear predictions about the relative response rates of the two groups can be derived. Indeed inspection of the response rates for the two groups suggests that they did not differ. This impression is confirmed by analysis of two sample sessions. On the first session the average response rate for the marked group of 161 did not differ significantly from the control average of 206, $t(12) = 0.94$, $p > .40$. The same was true for the 10th session where the rates for the marked and control groups were 290 and 314 respectively, $t(12) = 0.35$, $p > .50$.

The addition of the onset of the magazine light to the marker used in Experiment 5 was not successful in producing a marking effect in this experiment. If anything, the impact of this extra element was to detract from the efficiency of houselight and keylight offset alone. This outcome might have been due purely to chance. Alternatively it could have been caused by a reduction in the degree of change in illumination, a distractor effect, or the secondary reinforcer qualities of the magazine light.

3.5 GENERAL SUMMARY AND DISCUSSION

The work reported in this part of the thesis was directed at determining whether or not marking can be found using a visual discrimination task with pigeons in an operant chamber. It involved the manipulation of various parameters of the experimental situation in an attempt to find out whether a marking effect could be achieved under certain circumstances in this paradigm, or whether there was a fundamental limitation on the generality of the marking phenomenon. The results were not encouraging as in no case did following a choice response with a salient event significantly facilitate learning.

This discussion deals first with the outcomes of the various parametric manipulations which have been performed. It then goes on to examine the various hypotheses which have been proposed so far to account both for the failure to achieve a marking effect and for the occasionally superior performance of the control group. This section focuses primarily on Experiments 5 to 8. The next stage in the discussion considers evidence that marking can be found in spatial operant discriminations in pigeons - a discovery which suggests that the limits on the generality of the phenomenon indicated by the research reported here could be related to the visual nature of the task used. Finally, some conclusions about the implications of these experiments for marking are presented.

Three parameters have been varied within the eight experiments that have been presented so far. These are delay of reinforcement, delay-interval procedure, and type of marker. In the first four experiments delay of reinforcement was always 15 sec whereas in all subsequent experiments it was reduced to 6 sec. The reduction in delay did not seem to make much difference to performance. If the results of the whole-key-marker and control groups in Experiment 4 are compared with results for the marked and control groups in Experiment 5, then it can be seen that the level of learning was similar in both experiments (see Figures 8 and 10). In each case there was an increase of about 20 percentage points between the first and last training blocks.

A second set of manipulations involved the sequence of events that filled the delay interval. In some of the experiments choice responses were followed by the darkening of the non-chosen side of the response key. It was clearly established by Experiments 1 to 4 that the impact of this was to allow learning through perceptual focusing as a consequence of differential stimulus exposure. Another variation in procedure involved the maintenance or not of the discriminative stimuli until the end of the delay interval. When these stimuli were replaced with white illumination midway through the delay as in Experiment 5, the level of learning was generally much lower than when they remained on the response key until the end of the delay as in Experiment 7. Apart from this difference in delay-interval procedure, the marked and blackout-marker groups of Experiments 5 and 7 respectively were identical in design, as were the control groups of the two experiments. It can be seen from Figures 10 and 13 that for each pair of groups the increase in percent correct in Experiment 7 was approximately twice that in Experiment 5.

Several different markers were used both on and off the response key. The most damaging markers were found in Experiment 3, where choice responses were followed by a burst of white noise and houselight offset for 2 sec, and Experiment 6, where choices led to amber illumination of the key for 0.5 sec. In each instance the marked group did significantly worse than the control group. The most effective event to be used as a marker was the combination of houselight and keylight offset, both for 0.5 sec, in Experiment 5. The outcome in this case was a consistently higher percentage of correct choices by the marked group though the difference was not significant.

Several hypotheses have been proposed to account for the findings of these experiments. Before entering into a discussion of these hypotheses, an important point about the interpretation of the experimental work presented here should be mentioned. In selecting markers the intention was to choose a highly salient event of a kind which seemed likely to initiate a backward memory scan and thereby make a choice response more memorable. To achieve the required level of salience several elements were often combined to produce a compound marker. In retrospect this procedure can be seen to have made the analysis of results more difficult - it is possible that the effect on performance of each element of a marker was different. Thus there are problems in disentangling the impact of the various elements of markers in a way which would permit a more coherent assessment of the hypotheses presented. This problem renders the discussion which follows more tentative than would probably have been the case if single events rather than compound markers had been generally used.

Each of the hypotheses that has been proposed to account for some aspect or other of these experiments will now be considered and evaluated in the light of the pattern of results as a whole.

Interference hypothesis

One of the hypotheses is an interference hypothesis of the kind advanced by several investigators including Grant and Roberts (1976), Roberts (1980), and Wagner (1976, 1978). In a simple form the interference account suggests that events compete for processing time in a limited capacity short-term memory and that the impact of new events can be to displace preceding ones from memory, thus denying them the rehearsal necessary for learning. As applied to the marking work reported here, this interpretation implies that events during the delay, including the marker, may act to prevent rehearsal of information about the discriminative stimuli and the choice response, thereby damaging performance. In earlier discussions two versions of the interference hypothesis have been presented, one of which was general in scope and allowed that stimuli of all kinds could displace information from short-term memory with detrimental consequences. The other was more specific and proposed that the colour processing required in the red-green discrimination employed here would be particularly susceptible to disruption by colour stimuli occurring during the delay.

The general interference hypothesis was intially suggested as a possible explanation of failure to get marking in Experiment 5 and of the relatively low level of learning in situations where the red and green stimuli were replaced with white during the delay. This change in stimuli and the houselight-offset element of the marker were seen as potentially interfering events. Following the outcome of Experiment 6 where the marked group did significantly worse than the control group, the interference interpretation was extended to include the marker itself as an event that might displace information from short-term memory. The general interference hypothesis would predict that marked group performance in Experiments 5 and 8, and blackout-marker group learning in Experiment 7 should be worse than the performance of their respective control groups. This is because in each instance the marked group experienced greater stimulus change than the control group. In no case was this prediction confirmed and therefore the general pattern of results did not support the hypothesis.

Evidence from other researchers has mixed implications for a general interference hypothesis. In some cases additional stimuli have been found to interfere with performance, (see, for example, Grant and Roberts, 1976; Pearce and Hall, 1978). However, in those situations which are most similar to the experiments reported here no interference effect has been found. In work involving discrimination tasks with a salient event following both correct and incorrect choice responses (Fuller, 1981; Lieberman, personal communication, 1981; Lieberman et al., 1979; Thomas et al., 1983) facilitation, not interference, was observed. In particular, in experiments by Fuller (1981) and Lieberman (personal communication, 1981) involving spatial discrimination learning by pigeons using the same apparatus, the performance of the marked group was better than that of the control group.

The colour version of the interference hypothesis receives more support from these experiments. The amber-marker group in Experiment 6 performed significantly less well than the control group. The outcome of an attempted replication of these groups in Experiment 7 was unclear, but if the average results of the amber-marker subjects are replotted then they appear to do less well than either blackout-marker or control birds. These findings, especially when set against the general trend of results with other markers, suggest that colour markers may interfere with the specific colour processing needed for the solution of this particular discrimination. It should be noted that in none of the marking work by other researchers reported in this thesis was there a parallel similarity between the nature of the marker and the discrimination task (Fuller, 1981; Lieberman, personal communication, 1981; Lieberman et al., 1979; Thomas et al., 1983). What is not clear is whether this possible effect is one that is peculiar to colour stimuli, or whether it is more generally the case that when the discriminative stimuli and marker are drawn from the same stimulus class, then interference occurs. This question could be investigated using various combinations of stimuli drawn from two classes - for instance, colours and black-and-white patterns - as markers and discriminative stimuli using an approach similar to that of Grant and Roberts (1976). In a delayed matching-to-sample situation they varied the relationship between the sample stimulus and a stimulus interpolated during the interval between the sample and cmparison stimuli. They were thereby able to test whether interpolated stimuli drawn from the same category as the sample were more likely to interfere with matching than stimuli from a different category.

One way of evaluating the likely validity of the specific form of the interference hypothesis is to examine the evidence available from other, nonmarking, research. A similar line of argument to that proposed above has also recently been advanced by Mackintosh (1983, pp. 276-7) in a discussion of discrimination learning. In the context of an evaluation of the relationship between Wagner's (1978, 1981) rehearsal model and the effects of interpolated stimuli in delayed matching-to-sample tasks, Mackintosh proposes that the most satisfactory evidence of an effect of retroactive interference on retention comes from "studies that have shown that the amount of interference depends on the relationship between the interpolated event and the sample that is to be remembered" (Mackintosh, 1983, p.276). He cites work with monkeys by D'Amato in which performance in visual delayed matching-to-sample was impaired by the introduction of houselight between the sample stimulus and comparison stimuli. If the level of illumination, and consequently visual stimulation, was reduced during the delay interval then matching was facilitated (D'Amato and O'Neill, 1971). Also if auditory stimuli in the form of white noise or monkey vocalisations were inserted in the delay, no interference with matching was found (Worsham and D'Amato, 1973). As Mackintosh suggests, these studies imply that stimuli of a particular class are rehearsed in a relatively specific rehearsal mechanism. As a result, matching tasks with visual stimuli are peculiarly susceptible to disruption by other visual events. Evidence from studies of human short-term memory is consistent with this view (see, for example, Kroll, Parks, Parkinson, Bieber and Johnson, 1970; Peterson, Rawlings and Cohen, 1977; Salzberg, Parks, Kroll and Parkinson, 1971).

Further evidence that visual stimuli can interfere with performance is provided by studies of delayed matching-to-sample in pigeons. Grant and Roberts (1976), (whose work was referred to earlier), concluded from their research that degree of illumination during the delay interval was a significant source of damage to the matching performance of their pigeons. Other investigators (Cook, 1980;

Tranberg and Rilling, 1980) have also shown that changes in delay-interval illumination interfere with pigeon short-term memory in delayed matching-to-sample tasks.

Taken together, these studies of monkey and pigeon delayed matching-to-sample and human short-term memory support an interference model which allows that stimuli from a specific class are more likely to interfere with the rehearsal of stimuli from the same class than from a different class. At a theoretical level the results from other research areas are therefore in line with the potential implications of the experiments carried out for this thesis. However, at an empirical level the results of the two sets of experiments conflict on the question of how broadly defined the classes of stimuli are. On the one hand the marking experiments reported here suggest that interference operates at the level of colour stimuli. When visual stimuli were used as markers, such as houselight and keylight offset in Experiment 5, they did not produce a decrement in learning relative to the control group. On the other hand are the findings of Grant and Roberts (1976). In a series of experiments of the effects of interpolating stimuli between the sample stimulus and the test stimuli, they used stimuli from two classes - colours and black-and-white patterns. The objective was to test the hypothesis that "interpolated stimuli from either class would interfere more with delayed matching-to-sample of stimuli from its own class than those from the other class" (Grant and Roberts, 1976, p.2). No such effect was found. As was mentioned, they concluded that degree of illumination during the interval was the important factor.

In summary therefore, two problems arise when the categories of stimuli suggested by the two areas are compared. The work on delayed matching-to-sample and human short-term memory provides evidence of interference at a broad visual level which points to some kind of visual rehearsal mechanism. In contrast, the results of the marking experiments reported here suggest a more restricted interference effect at the colour as opposed to visual level. Not only is there a divergence between the two sets of results in terms of the breadth of categories within which interference is found, but the delayed matching-to-sample work by Grant and Roberts (1976) provides evidence against the notion of a specific colour rehearsal mechanism. Therefore non-marking research does not unequivocally support the kind of interference hypothesis suggested by the experiments conducted for this thesis. Further research would be useful, first, to establish whether there is definitely a colour interference effect in the paradigm used here; second, whether this effect is found with other classes of stimuli, and third, to examine which factors might be responsible for the apparent difference between delayed matching-to-sample and marking situations.

One result which does not fit the pattern suggested by the discussion of interference hypotheses above is the finding in Experiment 3 of slightly better performance in the control group than the marked group (see Figure 6). No learning was found in the marked group whose choice responses were followed by a combination of a burst of white noise and the offset of houselight, both for 2 sec. The control group showed an improvement in performance though it did significantly better than the marked group on only two sessions. Later experiments that also involved houselight offset, (though only for 0.5 sec), did not reveal a deficit in performance relative to controls, which suggests that the burst of white noise may have been responsible for the absence of learning in the marked group in Experiment 3.

Distractor hypothesis

A second hypothesis was raised in the discussion of these experiments which may be able to account for the anomalous result of Experiment 3. This distractor hypothesis rests on the assumptions that subjects benefit both from continued exposure to the discriminative stimuli and from repeated pecks to the chosen colour. It suggests that the impact of an off-the-key marker may be to distract the subject's attention away from the response key. The distractor interpretation as applied to Experiment 3 implies that the white-noise burst drew marked subjects' attention from the discriminative stimuli and thereby lowered their performance. A similar account was offered for a possible function of the houselight-offset element of the marker in Experiment 5, and of why the addition of magazine-light onset to the marker in Experiment 8 may have undermined the effectiveness of the other marker elements. The validity of the distractor hypothesis was considered at length in the discussion of Experiment 5 and can be questioned on three grounds. First, in Experiment 4 the whole-key-marker group, which received no exposure to the discriminative cues following choice responses, did no worse than the control group, an outcome which implies that the continued availability of cues does not help learning. Second, if the attention of marked subjects was distracted from the key, one might expect their response rates to be lower than that of control subjects. A comparison of response rates in Experiment 5 was made and although the rate in the marked group was lower, the difference between it and the control group was not significant. Third, the successful use of the white-noise and light-flash markers by Lieberman et al. (1979) and Thomas et al. (1983) was not consistent with a distractor view. There is therefore little support for a distractor hypothesis.

Adventitious learning

Two hypotheses were raised in the context of explaining why learning was so much better when the discriminative stimuli were maintained throughout the delay. One, which was based on a general intereference account, can now be discounted following the failure to find any support for it in the discussion earlier in this chapter. An alternative and more convincing interpretation of the improvement in performance found when discriminative stimuli are maintained is based on adventitious learning. It was found that birds continued to peck the response key during the delay following a choice. If it is assumed that there is an inertial tendency in their response patterns, then the subjects will receive occasional chance pairings of pecks to the correct colour with reward and to the incorrect colour with non-reward. Adventitious learning should therefore be possible. Unfortunately, the evidence of an inertial tendency needed to provide a firm basis for this interpretation is not available.

Startle hypothesis

One final hypothesis, which was suggested initially to explain the failure of the marked group in Experiment 6, no longer seems credible. This "startle effect" account proposed that the effect of following choices with an amber marker was to disrupt the inertial tendency and thereby remove the basis for adventitious reinforcement. Once again inertial tendency data may have helped in the assessment of this account, but in the light of Experiment 7 where a blackout-marker group learned as well as the control group, the startle view seems implausible. One would have expected the complete darkening of the chamber to be as likely to disrupt any inertial tendency as an amber flash. In any case it now seems that the results of Experiment 6 can be more adequately dealt with by a colour version of the interference hypothesis.

Conclusions

Discussion has now taken place of two interference hypotheses, a distractor hypothesis, an adventitious learning account, and a startle effect. The one interpretation to receive most, though not absolutely convincing, support from the data is the colour version of the interference hypothesis. The other accounts may in fact be correct with reference to the effects of particular elements of markers and could certainly be experimentally tested. However, the use of compound markers in some of this research and the absence of evidence about inertial response tendencies make it hard to reach more definite conclusions about the validity of these accounts.

What implications then does this series of experiments have for the generality of the marking phenomenon? This question is best answered in the light of other research which has been conducted using a spatial discrimination task. The experiments of Fuller (1981) and Lieberman (personal communication, 1981) have already been referred to. In both cases pigeons were tested using the same apparatus as in the research reported here and with essentially the same procedure. Discrete trials lasting 6 sec were employed with the response key illuminated white during the trial. In the marked groups, choice responses were followed by the darkening of the key for 0.75 sec. Learning was greater in the marked group. More recently, Davidson (1983) reported successful spatial marking with a rather different procedure but still using pigeons and the same split-key apparatus as in this work. (Full details of Davidson's design will be given in the next part of the thesis). Successes of this kind represent evidence that any limits to the generality of the marking phenomenon do not lie with the species used or with the experimental situation. Learning with pigeons in an operant chamber is facilitated when their choice responses are followed by a salient event if the discrimination is a spatial one. As the maze learning reported by Lieberman et al. (1979) and Thomas et al. (1983) also involved spatial tasks, it seems that there may be special difficulties with visual discrimination problems. A limit to marking's generality may lie with visual learning.

Why should marking in visual tasks be any more difficult to achieve than in spatial ones? One possibility is that visual discriminations themselves are more difficult than spatial problems. This could be because simultaneous visual discriminations necessarily involve two sets of cues, those relevant to reinforcement or non-reinforcement and those irrelevant to the solution of the problem. The red and green stimuli that the subject has to choose between for reward are presented in different positions - the left and right side of the response key - and this difference in position is irrelevant to the solution of the problem. The subjects might need to learn to attend to the relevant stimuli of the discrimination task (Mackintosh, 1983, pp. 244-5; Sutherland and Mackintosh, 1971). In contrast, the spatial discrimination problems used in the experiments by Fuller (1981), Lieberman (personal communication, 1981), and Davidson (1983) involved only one set of differential cues, those associated with position. In each of these experiments both sides of the response key were illuminated white. In other spatial marking experiments conducted using mazes (Lieberman et al., 1979; Thomas et al., 1983), additional, visual, cues were provided by painting the side arms (or the doors from the choice box to the side arms) black on one side and white on the other. However, the same colour was always associated with the same side. Therefore regardless of whether subjects initially coded the problem spatially or visually they still should have been able to learn about the problem. Thus in the spatial marking experiments either only one set of cues was present, or, if two sets were available, then both were relevant to the discrimination. Greater difficulty with the red-green discrimination used in the visual experiments reported here could have arisen because the subjects initially coded the problem in spatial terms and only later attended to colour. Evidence that pigeons find colour discriminations more difficult than spatial tasks is provided by Bullock and Bitterman (1962), and Schade and Bitterman (1966). Using a standard two-key apparatus, they found that even in an immediate reinforcement situation, (as opposed to the delayed reward used here), pigeons still made more errors on and took longer to solve a visual discrimination problem.

An analysis of why spatial marking in pigeons should succeed but visual marking fail, in terms of the relative ease of spatial and visual tasks, runs into an obvious problem given the results of this research. The difficulty is that although visual tasks may be harder than spatial ones, nonetheless the discrimination used here is capable of solution. In all the experiments, (even where adventitious reinforcement was not possible), learning did occur. The question therefore remains of why, given that the task could be solved, did marking not facilitate learning?

One way around this problem is suggested by considering those qualities that may be important in determining the effectiveness of a marker. Salience is one dimension, novelty has been proposed as another. A new, and therefore surprising, event is seen as more likely to initiate a memory search than one which has been previously experienced. This analysis suggests that a marker may be at its most effective during the early stages of training. If the additional assumption is made that the initial coding of the problem by the pigeons is in spatial terms, then it could be argued that the failure of marking in visual tasks arises because at the stage when the marker is most effective, subjects are attending to the wrong dimension. At a later point in training when the birds attend to the relevant, visual, cues the marker is no longer very effective. An account of this kind is capable of explaining the apparent difference between performance in spatial and visual tasks.

Unfortunately evidence from the response patterns of the subjects during the pretraining stages of the experiment does not support this interpretation. Many subjects do indeed show strong side preferences but there are also clear colour preferences at the very earliest stages of the experiments. Figures for the level of preference for the subsequently-designated correct colour during pretraining are given by way of illustration for two experiments. In Experiment 4, the average percentage of responses to the colour subsequently rewarded in training was 31%. The figure for Experiment 8 was 20%. Additional evidence that the subjects are indeed attending to colour is provided by the rapid initial improvement in performance found in several experiments. For example, in Experiment 6 both groups showed a rise of at least 12 percentage points in their level of correct responding between the first and second blocks, (see Figure 12). In Experiment 7, the results of which are shown in Figure 13, there was also a substantial improvement by all groups of more than 10 percentage points between the baseline provided by pretraining and the first block of delayed reward training.

It could be argued that although initial coding may be in visual terms, memory for colour is poor. However, this proposition runs into the problem that it does not differentiate between marked and control groups. Therefore it does not explain why marking, as opposed to learning, might be harder in visual tasks.

The discussion above has not identified any general feature of visual problems which could explain why it appears that marking might be possible in spatial but not visual discriminations. This suggests that the difficulty might lie with the particular procedures which have been adopted in this research programme. This conclusion is strengthened by the fact that the marking effect obtained using a spatial equivalent of the procedure (Fuller, 1981; Lieberman, personal communication, 1981) was not as powerful as in earlier maze experiments (Lieberman et al., 1979; Thomas et al., 1983). The final stage of the research therefore involved an investigation of visual marking using a quite different procedure - this work is described in the next part of the thesis.

What initial conclusions can be drawn about marking from the experiments reported so far? It does appear that there are limitations to the range of situations in which following choice responses with a salient or novel event will facilitate learning. Several events which intuitively should be highly salient for a pigeon, such as the complete darkening of the chamber, have not produced additional learning in the marked groups. An entirely novel event of the kind involved in an amber flash on the response key also failed to improve performance above that of the control group. Therefore it seems that salience or novelty is not necessarily enough to ensure that a marking effect will be revealed by facilitated performance. It may have been the case that marking occurred in these experiments in the sense that a backward scan through the memory store to identify a causal or predictive one took place, but if this did happen it was not reflected in improved performance. One problem that has been identified and which may have hampered the emergence of a marking effect, is the use of multiple events or compound markers in some experiments. The impact of the various elements on learning might have been in different directions, making it difficult to interpret some results clearly.

Another finding which has emerged is that events intended as markers can impair performance relative to a control group - a counter-marking effect. A colour interference hypothesis was proposed to account in part for this discovery. Further work on the relationship between the classes of stimuli used in a discrimination problem and as markers, and the consequences of different relationships would be profitable.

The next set of experiments was intended to unravel answers to some of the questions identified above.

CHAPTER 4

INVISIBLE-TRIAL EXPERIMENTS

4.1 GENERAL INTRODUCTION

There is little evidence from the experiments reported so far that supports the marking hypothesis. Following choice responses with a salient or novel event did not facilitate learning in the way that Lieberman et al. (1979) and Thomas et al. (1983) found. Several potential explanations for these results have been discussed. One possibility is that this failure could be attributed to either the species or apparatus used. The successful demonstration of marking with pigeons in operant chambers by Fuller (1981), Lieberman (personal communication, 1981), and Davidson (1983) did not support this view. A second explanation which has been proposed for the absence of a marking effect in these experiments is that the difficulty may lie with the visual nature of the discrimination task. All other investigations of marking have employed spatial discriminations. However, an examination of the theoretical and empirical evidence did not pinpoint any feature of visual problems that could clearly account for the failure to obtain marking. A third alternative is that the difficulty may lie with the particular procedure used in this research. This interpretation is encouraged by the fact that the marking effect achieved with a spatial version of the procedure was not powerful (Fuller, 1981; Lieberman, personal communication, 1981). The weakness of this effect led to the development of a new procedure for examining marking in pigeons in an operant chamber (Davidson, 1983). The effect obtained using Davidson's (1983) procedure was much stronger than in the spatial discrete-trial experiments, suggesting that it would be profitable to establish whether a similar improvement would be found with a visual version of Davidson's basic design. The next section of this discussion will describe Davidson's (1983) work and subsequent experiments by D. A. Lieberman (personal communications, June 1983, September 1983) using the same general design. The aim of this account is to provide a proper background to the visual version of the new procedure, and to allow a subsequent comparison of spatial and visual situations.

Experiments 1 to 8 and the work performed by Fuller (1981) and Lieberman (personal communication, 1981) all involved a discrete-trial procedure. The response key was darkened during the interval between the end of a trial and the next presentation of the discriminative stimuli. One problem with these experiments was that subjects in the control group did reasonably well, thus making it harder to demonstrate any improvement in learning produced by marking. Consequently in order to make the task more difficult, Davidson (1983) adopted what could be called an "invisible-trials" procedure. During delayed reward training trials were no longer signalled by the onset and offset of the discriminative stimuli. Instead the key remained illuminated for the entire duration of each session and trials were programmed at particular times by the experimenter. The first response to occur after a trial had begun was designated the choice response and initiated the delay interval.

Davidson (1983) used three groups of pigeons. In the choice-marked group, the first response of a trial produced a marker regardless of whether or not it had been to the correct side of the key. The marker involved the replacement of the white keylight with half-red, half-green key illumination for 1 sec immediately following the "choice" response. (The key was lit red and green not only to provide a salient marker but also to emphasize the two halves of the key and thereby make the discrimination easier.) After a further 6 sec the subjects were rewarded if their choice had been correct. In Davidson's second group - called the pseudo-marked group - subjects did not receive a marker on trials. Choice responses simply determined whether or not food was received 7 sec later and no change in key illumination occurred. However, 15 or 20 sec after each trial a pseudo-trial was programmed. The first response during this pseudo-trial by subjects in the pseudo-marked group produced a marker but had no other consequences. This group was introduced to control for arousal effects which might have been produced by the marker. A rise in the level of arousal may have increased attention and produced better learning through a non-marking process, (see the earlier discussion of the arousal hypothesis and the maze experiments). A final, control, group was also included which received no markers at any point in the experiment. The hope was that the new procedure would make the discrimination substantially harder for the control group to learn as they would receive no indication of which response was responsible for reinforcement.

Davidson found that performance in all three groups improved over the course of the experiment. The marking hypothesis predicted that learning should be best in the choice-marked group and this was borne out by the experimental results. The choice-marked group did significantly better than the pseudo-marked and control groups which did not differ from each other. The absence of any difference in percentage of correct choice responses between the pseudo-marked and control groups added to the evidence provided by Lieberman et al. (1979) and Thomas et al. (1983) against an arousal interpretation. However, it may be that arousal at the time of food delivery is the crucial factor and that the marker produces only a temporary increase in arousal. If these assumptions are correct, then an arousal hypothesis could also account for the obtained results since a marker which immediately preceded food, as in the choice-marked group, would be much more effective in facilitating learning. Thus the pseudo-marked group with its marker about midway between trials may not have been a good control for arousal interpretations of the function of markers. It did however control for the possibility that simply following responses occasionally with a salient event might somehow enhance learning by, for instance, drawing attention to the response.

A subsequent experiment using the invisible-trials procedure by D. A. Lieberman (personal communication, June 1983) incorporated a more effective control for an arousal view of marking. The design of the experiment was the same as Davidson's (1983) except that the pseudo-marked group was replaced with a delayed-marker group. In this group the presentation of the marker was delayed until 3 sec after the subject had made its choice response. A design of this kind was seen as providing a control for interpretations based on cycles of arousal.

Lieberman found that a number of subjects stopped responding during training and that the majority of these were control group subjects. As a consequence a percentage correct analysis of the birds' performance overestimated the level of learning in the control group. Accordingly a different measure of learning was used to analyse the results - the number of correct choice responses made in each session. The marker and delayed-marker groups both showed significantly better learning than the control group but did not differ from one another. Lieberman hypothesized that the absence of a difference between the performances of the two marked groups might have arisen from adventitious marking in the delayed-marker group as a result of an inertial response tendency. Following a choice response, subjects might have continued to peck to the same side of the key, thereby producing chance pairings of a response with the delayed marker. In order to test this hypothesis Lieberman examined the responding of the delayed-marker group over the final three sessions of the experiment when its behaviour was stable. He found that the last response to precede a marker on rewarded trials was correct 81% of the time. This evidence supports the view that learning might have occurred in the delayed-marker group because of adventitious marking.

The third experiment using the invisible-trials procedure (D. A. Lieberman, personal communication, September 1983) was designed to provide a more sensitive test of the arousal and memorial accounts of the way markers operate. Once again there were three groups of subjects but the delay of reward was increased to 10 sec and the trial procedure was changed in an important way for the marked groups. In the marked groups responses were no longer immediately marked. Instead markers could only be received during a "window" which opened 2 sec after the choice response and closed 7 sec later, that is 1 sec before the end of the delay. In one marked group - called marked-same - if the choice response had been to the correct side of the key, then the first correct response during the window was marked. Similarly, if the choice response had been incorrect, then the first incorrect response during the window was marked. Thus although the first response of a trial determined its outcome, marking did not occur until at least 2 sec, and at most 9 sec later. The reason for

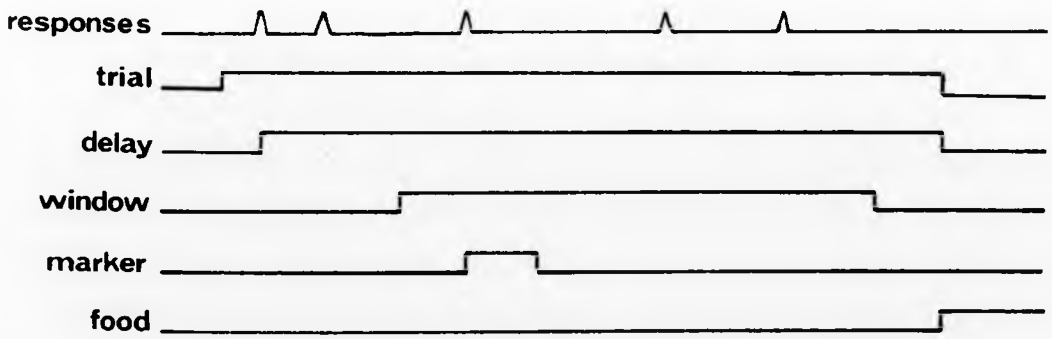

Figure 17

The "window" design used in D.A. Lieberman (personal communication, September 1983) and Experiments 9 to 11.

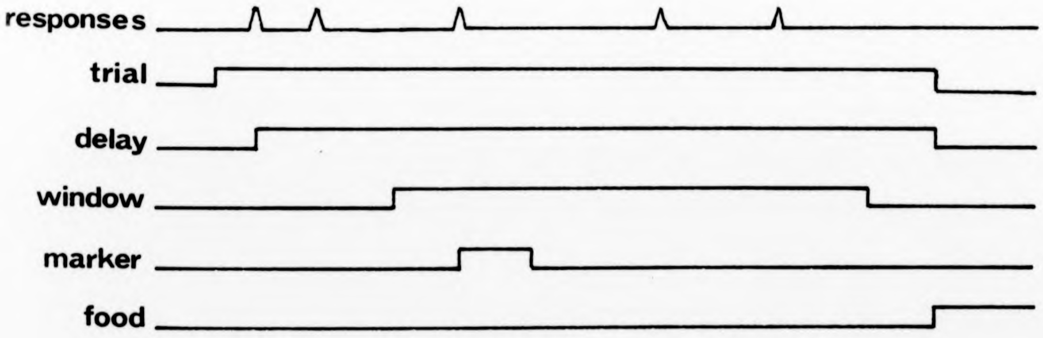

Figure 17

The "window" design used in D.A. Lieberman (personal communication, September 1983) and Experiments 9 to 11.

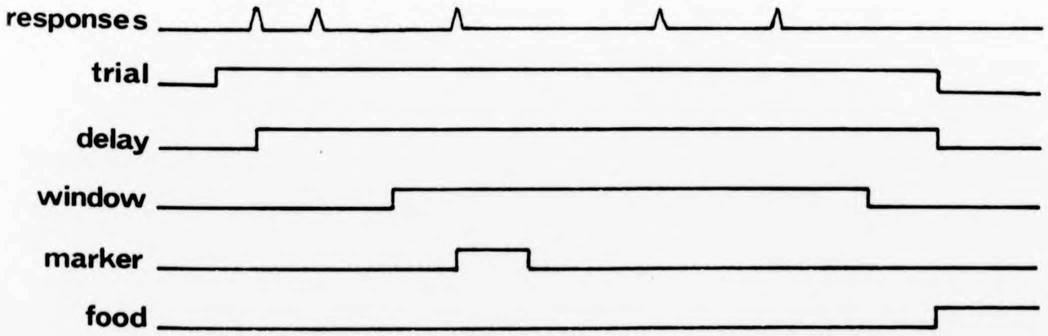

Figure 17

The "window" design used in D.A. Lieberman (personal communication, September 1983) and Experiments 9 to 11.

the adoption of this procedure as opposed to the more simple approach used in earlier experiments arose from the design of the second marked group. In this group - called marked-different - the choice response determined the trial outcome in the usual way. However, the first response to be marked during the window was not the same as the choice response. Rather, if a correct choice had been made, then the first incorrect response during the window was marked and vice versa. Thus the marker contingency was reversed for the marked-different group. The need for the postponement of marking (rather than simply marking the first response to occur after the choice that was the same as the choice in the marked-same group and different in the marked-different group), arose from the likelihood that subjects possessed an inertial response tendency. If they did, then it was probable that the marked-same subjects, (where the response to be marked was the same as the choice response), would receive a marker sooner than the marked-different subjects, (where the response to be marked was different from the choice response). As a result there would be a dissimilarity between the two marked groups not only in terms of the relationship between choice and marked response, but also in the delay between marker and reinforcement. This confounding of factors would make interpretation more difficult and therefore the 2-sec postponement of marking was introduced to allow time for the effect of any inertial tendency to be weakened and ensure that both groups received markers at about the same point in the trial. (The similarity of marker latencies during the first few sessions of training pointed to the success of this approach.) As in the experiments by Davidson (1983) and Lieberman (personal communication, June, 1983), the marker was a 1-sec change in key illumination from white to red and green. In the control group for the experiment, the choice response simply determined whether or not reinforcement was provided subsequently.

This experiment by Lieberman provided a very sensitive test of an arousal interpretation of the way markers operate. If the arousal view was correct, then equal learning would have been expected in the marked-same and marked-different groups. On the other hand if the memorial account of marking proposed by Lieberman et al. (1979) and Thomas et al. (1983) was accurate, then only the marked-same group should benefit from the marker. Only in that group would the marker initiate a backward scan through memory that would identify the correct response on rewarded trials and make it more likely to be recalled when reward was subsequently received. In the marked-different group the subjects were likely to attend to a response to the incorrect side as a result of the marker. If anything performance in this group should have been worse than in the other two as a result of the birds erroneously learning that incorrect responses led to food.

The outcome of Lieberman's (personal communication, September 1983) experiment provided convincing support for the memorial account of marking. However, for a second time, the percentage correct measure of learning was rendered unreliable because different numbers of subjects in each group dropped out during training. The number quitting in the control group was highest with the result that percentage data for that group were variable and generally overestimated the level of performance by its subjects. The outcome of the experiment was therefore analysed in terms of number of correct and incorrect responses, each measure being based on all the responses in a session except those during the delay interval following choices.

Responses occurring during the delay were excluded in order to eliminate any responses directly elicited by the marker.

The number of correct responses in the marked-same group was significantly higher than in the other groups and doubled over the course of training. There was also an increase in correct responding in the marked-different group while the level in the control group remained fairly stable throughout the experiment. The pattern for number of incorrect responses was also in line with predictions derived from the marking hypothesis. The level in the marked-different group showed a slight increase, whereas in the other two groups there was a substantial decline. Towards the end of training the rate in the marked-same group was somewhat above that in the control group.

The three experiments described above (Davidson, 1983; Lieberman, personal communications, June 1983, September 1983) provide a clear indication that marking is possible with pigeons in an operant chamber. The strength of the marking effect achieved using the invisible-trial procedure suggests that the more marginal results of spatial work by Fuller (1981) and Lieberman (personal communication, 1981), and the difficulties encountered in Experiments 1 to 8 may have arisen from the discrete-trial procedure employed. Taken together, these results imply that it would be valuable to examine marking with a visual discrimination using the invisible-trial procedure.

4.2 EXPERIMENT 9

In this experiment a visual version of the invisible-trial procedure employed by Lieberman (personal communication, September 1983) was adopted. The discrimination problem used was the same red-green one as in Experiments 1 to 8.

The primary objective of the experiment was to determine whether the difficulties in obtaining marking in earlier experiments were related to the visual nature of the discrimination or to the procedure used. If no marking effect was found in Experiment 9, then it would suggest that there are problems peculiar to marking in visual tasks. On the other hand, if a marking effect did occur it would imply that the difficulties in Experiments 1 to 8 arose from the procedures used.

Success in demonstrating a marking effect in this experiment would be helpful in several ways. One, it would provide a visual replication of Lieberman's spatial experiment. Two, it would involve a further test of the memorial and arousal accounts of marking. Three, identification of those features of the discrete-trial procedure that led to the absence of a marking effect would be facilitated. Finally, a paradigm for further exploration of visual marking would be available.

As in Lieberman (personal communication, September 1983) three groups were used in Experiment 9. In one of the marked groups, marked-same, the response that was marked was the same as the choice response. In the other marked group, marked-different, it was different. A new marker was introduced in this experiment - the response key illumination changed from red and green to white for 1 sec. As in Lieberman's experiment markers could only be received during a window in the delay. In the control group the choice response simply determined whether reinforcement was received.

The marking hypothesis would predict that the marked-same group should show the highest percentage of correct responses and the marked-different group the least. The latter result would be expected since a backward scan through memory following reinforcement should be likely to identify the incorrect response for subjects in the marked-different group. The effect on the performance of the control group in percentage terms was less clear in advance. It was not known whether several control subjects would extinguish in the way that happened in the experiment by Lieberman (personal communication, September 1983). If the drop-out rate did differ across groups, then percent correct would become an unreliable measure of comparative performance.

If the number of correct responses rather than the percentage is considered, then the marked-same group should show an increase in level of responding above that of either the marked-different or control groups. On number of incorrect responses, the marking hypothesis suggests that the marked-different group should do better than the other two groups.

This experiment was also seen as an opportunity to examine the role of inertial tendency in learning, and the effect of the marker upon inertial tendency. In several earlier experiments it had been suggested that an inertial tendency in the subjects' responding might have been the mechanism by which learning occurred, particularly in control groups. Inertial tendency might have promoted learning either as a result of adventitious marking in situations where the response key illumination changed during the interval, (for example, Experiment 5), or as a consequence of adventitious reinforcement when the discriminative stimuli were maintained until the end of the delay (for example, Experiment 6). It had also been suggested in Experiment 6 that marked subjects might have failed to benefit from adventitious reinforcement because the marker disrupted an inertial tendency - a "startle effect" - although this possibility was subsequently seen as unlikely. Unfortunately an error in data collection had made it impossible to evaluate directly whether or not an inertial tendency was present in birds' pecking in these earlier experiments. However, analysis of Davidson's (1983) data had revealed that an inertial tendency was present in all her groups.

In the present experiment determining the presence or absence of an inertial tendency was seen as valuable for three reasons. First, to confirm Davidson's findings; second, to provide indirect but useful evidence about the processes that might have operated in earlier experiments; and third, to obtain information which could help in the interpretation of the results of this experiment.

In order to measure inertial tendency two probabilities were calculated. One was the probability that the last response[1] to be made during a trial was correct given that the choice response on that trial had been correct. The second was the probability that the last response on a trial was correct given that the choice response on that trial had been incorrect. The difference between these scores was a measure of the subject's inertial tendency; if the final response on a trial was not influenced by the initial choice response,

---

[1] It was decided to base the measure upon the last response because if inertial tendency was to be capable of producing adventitious learning of the kind provisionally ascribed to it in earlier experiments, then it would have to be sufficiently strong to endure until the final response of the delay. The relationship between the choice and last response is therefore the one of greatest interest.

then the probability of correct given correct would be the same as the probability of correct given incorrect, and the difference between the scores would be zero.

A final aim of this experiment was to assess whether the marker had an impact on any inertial tendency. In order to be able to do this within groups as well as between groups, data were also collected for a "pseudotrial" which occurred half-way between real trials. No markers or reinforcements were programmed during pseudotrials. It was thus possible to compare a measure of inertial tendency during trials - when markers were received by marked subjects - with a measure from periods when no marker had been presented and thereby assess to what degree markers affected behaviour. If the marker was disruptive, then the level of inertial tendency during pseudotrials should be greater than during trials.

Method

Subjects

The subjects were 24 adult Rock pigeons which were experimentally naive at the start of the experiment. They were housed and maintained under the same conditions as in previous experiments.

Apparatus

In addition to the two operant chambers used in earlier experiments, another two BRS/LVE operant chambers were used in this experiment. They were identical in design to the original boxes. In all boxes a new mechanism for registering pecks to the response key was introduced.

This mechanism was more stable and reliable, and required less fine-tuning to ensure that its sensitivity did not vary.

White noise at about 65 dB (SPL), (measured 7 cm in front of the centre response key), was continuously supplied to each chamber for the duration of each session. The loudspeakers were mounted on the ceilings of the chambers directly behind the centre of the front panel. The houselight in each chamber was illuminated for the duration of every session, except in magazine training when it remained off until the subject had received its first reinforcement.

The electromechanical and solid state control equipment used in previous experiments was replaced with two Apple 2 computers which were located in a different room from both the home cages and operant chambers.

Procedure

Magazine training and shaping. The magazine training given in this experiment differed from that provided in Experiments 1 to 8. Before the start of the session the food magazine was raised and filled with grain, and the response key was illuminated with red and green lights. The subject was then placed in the chamber. After it had approached the magazine and fed for approximately 10 sec, the magazine was lowered and then immediately raised again. Thereafter magazine deliveries were progressively shortened until the bird was eating during a 3-sec magazine presentation. During this process, the intervals between presentations of the magazine were progressively lengthened.

Once the subject had learned to obtain food from the magazine, manually controlled shaping began with the aim of training it to peck at the response key. When the subject had received 10 consecutive reinforcements for pecks to the key, the reinforcement schedule was changed from continuous to random interval.[1] The initial probability of reinforcement per second was set at 0.2 and this was progressively reduced to 0.1 in steps of 0.01. A reduction in the probability of reinforcement occurred on every other reinforcement if the time since the last response or reinforcement was less than 3 sec. In order to establish responding and avoid extinction the equipment was programmed to provide increases in the probability of reinforcement if the subject's response rate fell beyond a certain level. Whenever the latency of a reinforced response was greater than 5 sec, then the probability of reinforcement was increased by 0.01. Once the subject reached a probability of reinforcement per second of 0.1, this stage of training was completed. The duration of each session was therefore variable.

Three subjects completed the magazine training and shaping stages within a single session, 12 subjects took two sessions, and the remainder three or four.

Concurrent variable-interval training. All subjects then received four daily sessions of concurrent VI training, designed to ensure high rates of responding to both sides of the key. During the first session reinforcement was simultaneously available on a VI 40-sec schedule on each half of the response key. For all but two subjects this session was followed by two sessions of concurrent VI 60-sec, and a final session of concurrent VI 90-sec. For one subject the

---

[1] Zeiler (1977, p.202) defines random interval schedules as "schedules that provide reinforcer presentation after irregular time periods ... but the precise sequences are not prespecified. Instead, each time period ... is equally eligible for reinforcement according to some probability". The delivery of reinforcement is, of course, response dependent.

second session was also concurrent VI 40-sec and another subject received five concurrent VI sessions including three of concurrent VI 60-sec. Each session lasted 30 min.

In Lieberman's (personal communication, September 1983) study, response key illumination was the same during concurrent VI training as during delayed reward training (except, of course, for presentations of markers during the latter stage). In order to ensure the same level of similarity between concurrent VI and delayed reward training in terms of cue availability in this experiment, the red and green stimuli were programmed to occur on both sides of the key during concurrent VI training. Changes from red-green to green-red, and vice versa, were scheduled to take place at the end of food reinforcement, and with a probability on any particular reinforcement of 0.5. Alternation of colours could occur at this point without it being likely that the subject was attending to the response key. In this way subjects' exposure to stimulus change before the training stage could be minimised.

Delayed reward training. The birds were divided into three groups of eight subjects. The groups were matched for level of colour preference, level of side preference, and response rate on the basis of behaviour in the final session of concurrent VI training. The running order of subjects and their allocation to each of the operant chambers was balanced across groups. Subjects were rewarded for choices to the colour that they had responded to least on the final concurrent VI session. Thosebirds which had shown no preference were rewarded in such a way as to maximise the matching of the groups.

For all subjects the response key was illuminated throughout the session and the same invisible-trial procedure as in Lieberman (personal communication, September 1983) was used. The first response after a trial began was designated the choice response. If it was to the correct colour, then 3-sec access to food was made available after a 10-sec delay. If it was to the incorrect colour, then the subject received no reinforcement. The red and green stimuli were alternated between sides of the key in the same way as during concurrent VI training.

For marked subjects the same window procedure as in Lieberman (personal communication, September 1983) was used, (see Figure 17). Markers could be received only during a period that started 2 sec after the choice response and finished 7 sec later. In the marked-same group the first response during this period that was the same as the choice response, was marked. In the marked-different group the first response that differed from the choice response was marked. In both groups the marker was the replacement of the red and green key illumination with white for 1 sec immediately following the marked response. In the control group no markers were given.

In all groups a pseudotrial was programmed to occur 10 sec after the end of a trial. The first response after the start of a pseudotrial initiated a 10-sec period. During this period responses had no programmed consequences and did not produced either markers or reinforcement. The only purpose of pseudotrials was to provide data for periods of the same length as trials but without marker presentations. This data could then be compared with behaviour during trials, allowing an assessment of the effect of markers upon pecking. The interval between the end of a pseudotrial and the start of the following trial was 10 sec.

All subjects received 20 sessions of delayed reward training, each session lasting 30 min. The number of trials in a session, up to a maximum possible of 45, was therefore dependent upon the subject's response rate.

Inertial response tendency was measured in the following way. Two probabilities were calculated. One was the probability that the last response to be made during a trial was correct given that the choice response had been correct. The second was the probability that the last response on a trial was correct given that the choice response on that trial had been incorrect. The second probability was then subtracted from the first and the product multiplied by 100 to give a measure of inertial tendency. Thus, a score of 100 would indicate a complete inertial tendency such that the final response on a trial was always completely determined by the initial response on that trial. If there was no inertial tendency and each response was completely independent of the preceding response, then the two probabilities would be equal giving a score of zero. The stronger the tendency, the higher the figure up to a maximum of 100.

In order to assess fully the impact of the marker upon any tendency to repeat responses, the inertial tendency for pseudotrials was also calculated. A measure of the effect of the marker was then obtained by subtracting the inertial tendency score for pseudotrials from the inertial tendency score for trials. If the marker disrupted inertial tendency on trials, then one would expect scores for the marker effect of less than zero. In the control group, the marker effect measure should be about zero as the conditions during both pseudotrials and trials were the same.

Results

Two subjects, one in each of the marked groups, made no responses after the first two sessions. As these subjects did not come into effective contact with the contingencies of the experiment, their data were excluded from all analyses.

Percent correct

Figure 18 shows the percentage of correct responses averaged over blocks of two sessions for each group in the experiment. This measure, which was based on choice and intertrial interval responses, showed no improvement in all groups over the course of the experiment from similar starting points in the first block. The highest level of learning was found in the marked-same group with the marked-different group doing least well. A two-way analysis of variance revealed a significant difference between groups, $F(2, 19) = 4.25$, $p < .05$, and a significant increase in percent correct over blocks, $F(9, 171) = 9.22$, $p < .001$. The interaction between blocks and groups was not significant, $F(18, 171) = 0.77$, $p > .70$. Subsequent Newman-Keuls comparisons (Winer, 1971, pp.528-529) between groups showed that the marked-same group performed significantly better than the marked-different group ($p < .05$). No other comparisons were significant.

Number of correct responses

The numbers of correct responses averaged over blocks of two sessions for each group are presented in Figure 19. This measure was based on all correct responses except those during the delay following a choice response. A very substantial increase in the number of correct responses by the marked-same group is apparent. In contrast, the level of correct responding rose only marginally in the marked-different group and dropped in the control. A two-way

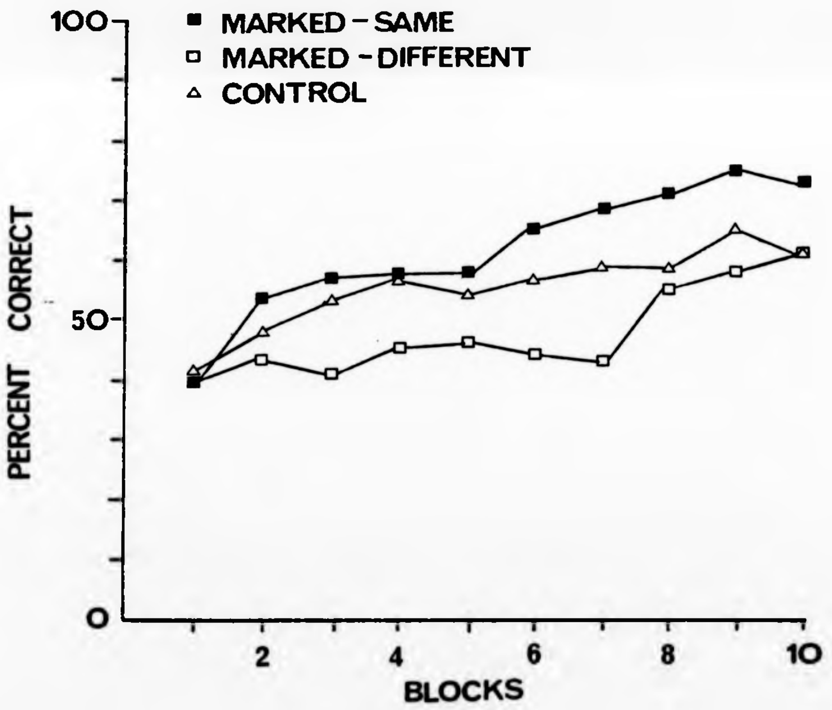

Figure 18

Mean percentage of correct responses over blocks of two sessions for each group in Experiment 9.

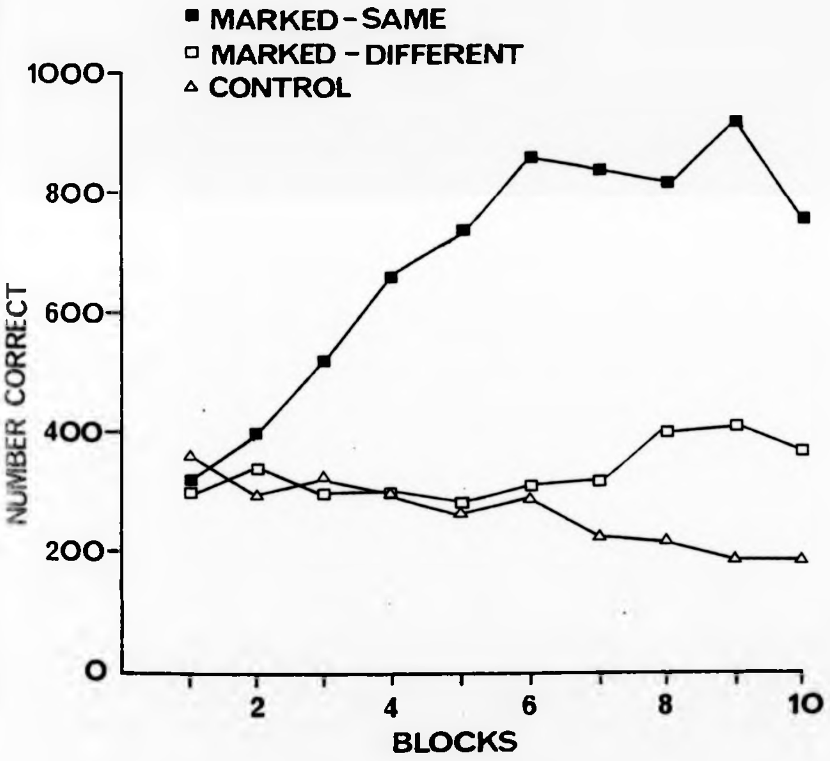

Figure 19

Mean number of correct responses over blocks of two sessions for each group in Experiment 9.

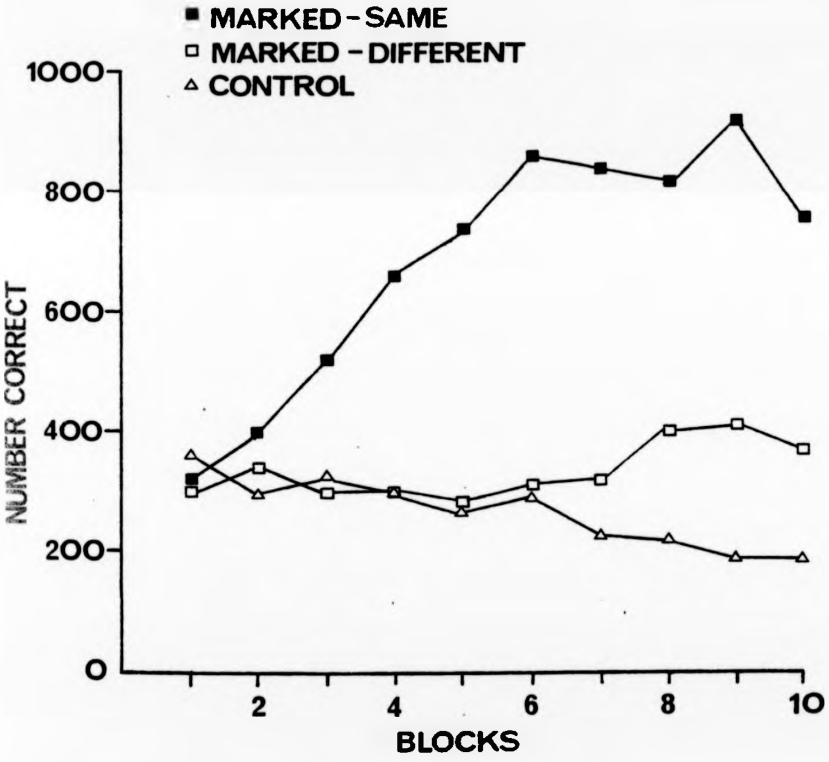

Figure 19

Mean number of correct responses over blocks of two sessions for each group in Experiment 9.

analysis of variance for the factors of experimental treatments and blocks of two sessions confirmed that the number of correct responses differed between groups, $F(2, 19) = 7.09$, $p < .01$. There was no change in overall number of correct responses over the course of the experiment, $F(9, 171) = 1.53$, $p > .10$, but the Blocks × Groups interaction was significant, $F(18, 171) = 2.74$, $p < .001$. Tests of simple main effects (Winer, 1971, pp.529-533) indicated that the groups differed on Blocks 4 to 10 inclusively, $Fs(2, 19) > 4.00$, $ps < .05$. Subsequent Newman-Keuls tests revealed that the marked-same group responded at a higher rate than both the marked-different and control groups on Blocks 4 to 10 inclusive, (all $ps < .01$ except on Block 4 where the difference between the two marked groups was significant only at the $p < .05$ level). The marked-different and control groups did not differ significantly from each other on any block.

Number of incorrect responses

The numbers of incorrect responses for each group averaged over blocks of two sessions are shown in Figure 20. As in previous measures, the only responses to be excluded from the analysis are those which occurred during the delay interval following choices. All groups showed a decline in the number of incorrect responses over the course of delayed reward training with the control group falling most. A two-way analysis of variance confirmed that the level of incorrect responding dropped significantly over blocks, $F(9, 171) = 10.10$, $p < .001$. The differences in the performance of the groups were not significant, $F(2, 19) = 1.67$, $p > .20$, and the Blocks × Groups interaction was also insignificant, $F(18, 171) = 1.36$, $p > .15$.[1]

---

[1] The absence of any statistically significant difference in incorrect responding between the two marked groups on the final block makes it unlikely that the difference in correct responding is due to the marker producing a general increase in responding in the marked-same group.

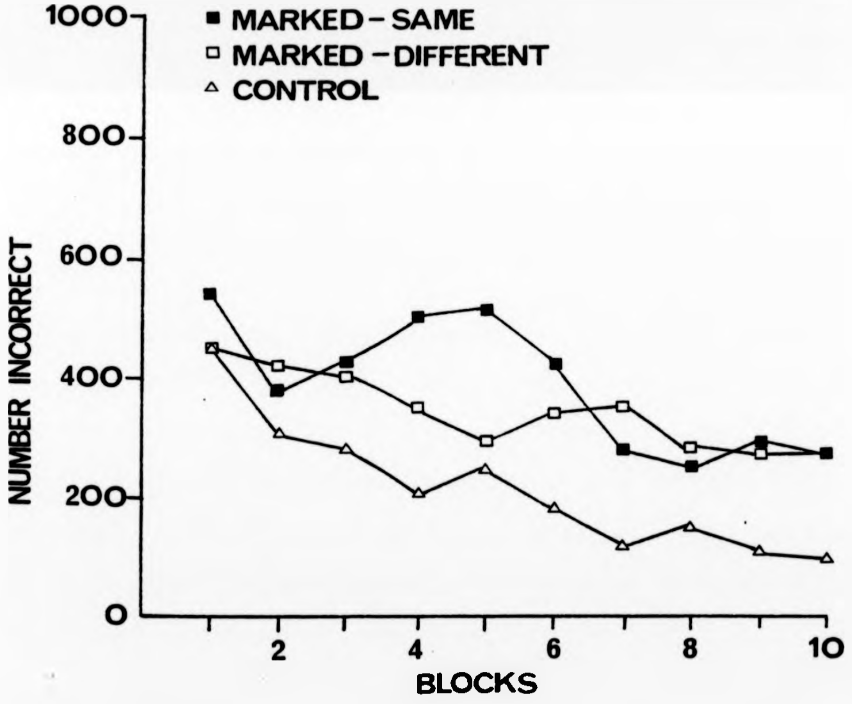

Figure 20

Mean number of incorrect responses over blocks of two sessions for each group in Experiment 9.

Inertial response tendency

The data were also analysed to give the measure of inertial tendency described in the procedure for this experiment. The level of inertial tendency for each of the groups averaged over blocks of four sessions is shown in Figure 21. There was no consistent difference between the groups nor was there any obvious change in the level of inertial tendency over the course of the experiment. These impressions were confirmed by a two-way analysis of variance that showed no effect of groups, $F(2, 19) = 0.40$, $p > .65$, or blocks, $F(4, 76) = 0.50$, $p > .70$. The interaction between groups and blocks was also insignificant, $F(8, 76) = 0.79$, $p > .60$. In order to determine whether or not there was a significant inertial tendency, the figures for each subject were averaged over delayed reward training and compared with zero using a one-sample *t*-test. The result was significant, $t(21) = 8.15$, $p < .0001$, showing that subjects did tend to repeat their most recent response at a level above chance.

Effect of marker upon inertial tendency

The effect of the marker upon inertial tendency was assessed in two ways. The first, which was based on the analysis of inertial tendency during trials reported above, showed no difference between the groups. The presentation of the marker in the marked groups did not disrupt, (or strengthen), the inertial tendency in those groups relative to that found in the control group.

The second assessment of the marker's effect involved a comparison of inertial tendency during pseudotrials (midway between real trials) with inertial tendency during trials. Figure 22 presents the data for the marker effect averaged over blocks of four sessions for each group. These data were subjected to a two-way analysis of variance

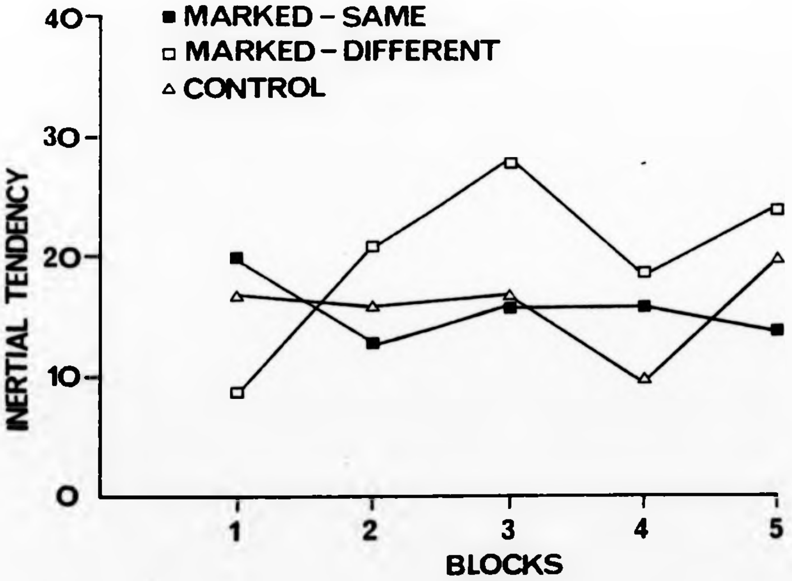

Figure 21

Mean inertial tendency over blocks of four sessions for each group in Experiment 9.

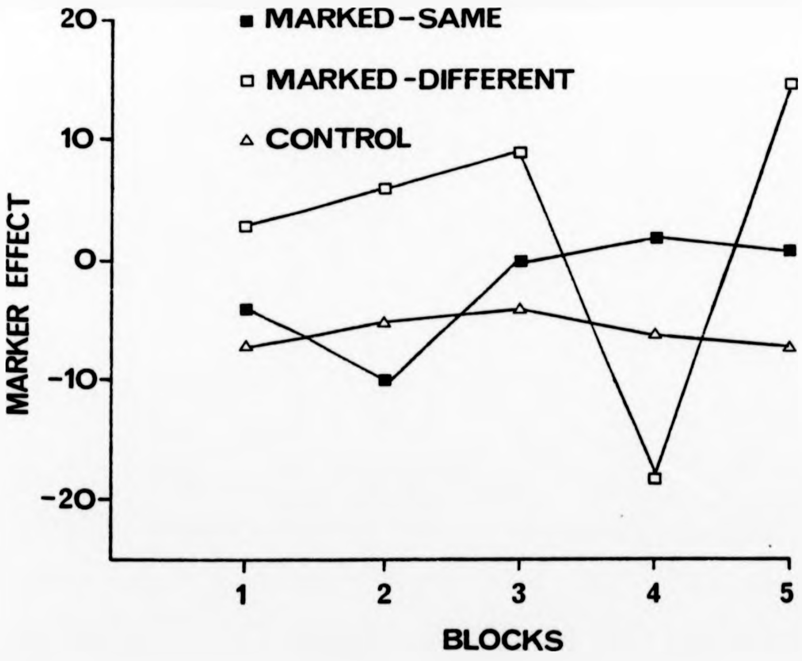

Figure 22

Mean effect of marker over blocks of four sessions for each group in Experiment 9.

which revealed no change in the effect of marker on inertial tendency over blocks, $F(4, 76) = 0.91$, $p > .45$. The interaction between groups and blocks was also insignificant, $F(8, 76) = 1.21$, $p > .30$. However, Figure 22 suggests that the level in the control group might have been lower than in the marked-different group and the presence of a significant difference between groups was confirmed by the analysis of variance, $F(2, 19) = 4.59$, $p < .05$. Subsequent Newman-Keuls comparisons showed that the figure for the marked-different group was significantly greater than that for the control group. $p < .05$. No other comparisons were significant. A one-sample t-test was then carried out to determine whether the performance of the control group differed from the expected value of zero and the test confirmed that, rather surprisingly, it nearly did, $t(7) = -2.24$, $p = .06$. In neither of the marked groups was the difference between the expected and obtained value for the effect of the marker significant. For the marked-same group $t(6) = -1.77$, $p > .10$, and for the marked-different group $t(6) = 1.53$, $p > .15$.

Discussion

This experiment provided much clearer support for the marking hypothesis than any of the previous ones conducted for this thesis. The results for percent correct were in line with predictions derived from the hypothesis. The marked-same group performed best with a level of responding that was higher than the control group's and significantly better than the marked-different group's. It showed an increase of 30 percentage points between the first and last blocks of delayed reward training compared with a rise of only 20 points for each of the other groups. The result of the marked-different group is interesting as it shows that following an incorrect response with a marker on rewarded trials can depress performance below that of a control group, (although the difference was not significant). It should be noted that the differences between the marked groups cannot be attributed to a difference in the timing of the marker. The window procedure was successful in ensuring that the latency of the marker following correct choices was very similar for each group. Averaged over the first five sessions, the marker latency in the marked-same group was 2.52 sec and in the marked-different group, 2.60 sec. (The subsequent divergence between the groups, with the average latency for marked-same subjects shortening, can be attributed to an increased response rate in those subjects.)

The number of correct responses also provided very clear evidence of a marking effect. The substantially higher rate of responding in the marked-same group clearly indicated that following a correct response on rewarded trials with a marker facilitated learning about that response.[1]

The pattern of results represented by number of incorrect responses in each group, shown in Figure 20, fits slightly less well with the predictions of the marking hypothesis. As expected the level of incorrect responding was greater in the marked-different group than in the control group though the difference was not significant. However the rate in the marked-same group was as high as in the marked-different group even during the final stages of the experiment. This result is slightly surprising as one might have expected a higher rate in the marked-different group because in that group a backward

---

[1] See footnote 1 on page 178.

scan through memory following reward should identify an incorrect response as the causal or predictive one. This result will be returned to and examined in greater detail in the discussion of Experiment 11.

As well as providing evidence to support the marking hypothesis, the results of Experiment 9 also make it difficult to sustain two alternative accounts of how the marked-same group did better than the control group. One such account, based on arousal, has been introduced in discussions of earlier experiments by Davidson (1983), Lieberman (personal communications, June 1983, September 1983), Lieberman et al. (1979), and Thomas et al. (1983). If the arousal view was correct it would predict that learning in the two marked groups should be the same. This is manifestly not the case and Experiment 9 therefore confirmed that the arousal hypothesis is as inappropriate in a visual situation as it has been shown to be in spatial tasks. In neither situation is it a sufficient account of the results.

A second interpretation of the relative performance of the marked-same and control groups could be offered in terms of generalisation decrement. As Mackintosh (1974, pp.406-407) stated, "a response which was established under one set of conditions will, by generalisation decrement, be less likely to occur when those conditions are changed". Applying this principle to the current experiment, it is at least possible that the conditions of pretraining established some unspecified behaviour that would retard learning during the subsequent delayed reward training. If so, there could then be an advantage to the marked-same group compared with the control group when it came to learning the discrimination as the introduction of the marker would scan through memory following reward should identify an incorrect response as the causal or predictive one. This result will be returned to and examined in greater detail in the discussion of Experiment 11.

As well as providing evidence to support the marking hypothesis, the results of Experiment 9 also make it difficult to sustain two alternative accounts of how the marked-same group did better than the control group. One such account, based on arousal, has been introduced in discussions of earlier experiments by Davidson (1983), Lieberman (personal communications, June 1983, September 1983), Lieberman et al. (1979), and Thomas et al. (1983). If the arousal view was correct it would predict that learning in the two marked groups should be the same. This is manifestly not the case and Experiment 9 therefore confirmed that the arousal hypothesis is as inappropriate in a visual situation as it has been shown to be in spatial tasks. In neither situation is it a sufficient account of the results.

A second interpretation of the relative performance of the marked-same and control groups could be offered in terms of generalisation decrement. As Mackintosh (1974, pp.406-407) stated, "a response which was established under one set of conditions will, by generalisation decrement, be less likely to occur when those conditions are changed". Applying this principle to the current experiment, it is at least possible that the conditions of pretraining established some unspecified behaviour that would retard learning during the subsequent delayed reward training. If so, there could then be an advantage to the marked-same group compared with the control group when it came to learning the discrimination as the introduction of the marker would represent a change in conditions between the two stages of the experiment. Behaviour established by pretraining would consequently be less likely to persist in the marked-same group. Two features of behaviour that might develop in pretraining and which would interfere with discrimination learning can be identified. The first is that training with concurrent VI schedules, particularly of the same value, can encourage switching behaviour (Catania, 1962; Herrnstein, 1961; Skinner, 1950). Alternation of responses would make it more difficult for subjects to learn through adventitious reinforcement. Therefore if generalisation decrement reduced switching behaviour in the marked-same subjects they would be able to benefit more from adventitious reinforcement than control subjects. Direct evidence against this version of generalisation decrement is provided by inertial tendency data for the marked-same and control group. Although on the first block of delayed reward training the inertial tendency was higher in the marked-same group, 19.9 as against 17.1 in the control group, the difference was slight.

A second, and more important feature of pretraining is that it might have reduced attention to colour because colour was not relevant to reinforcement at that stage. If learning that colour was irrelevant was undermined by generalisation decrement arising from the introduction of the marker, then this could help account for the superiority of the marked-same group over the control group.

Once again however, the inclusion of marked-different subjects provide an important comparison which indicates that a marking effect is responsible for the difference between the results of the marked-same and control groups. As with the arousal hypothesis, the generalisation decrement account predicts that both the marked-same and marked-different subjects should benefit from the introduction of the marker. The poor performance of the marked-different group, which only rose above 50% correct towards the end of training is therefore at variance with this account.

This discussion has so far examined the results of Experiment 9 in terms of predictions derived from the marking, arousal, and generalisation decrement hypotheses. Although the results provide clear support for the marking hypothesis, the learning found in both the control and marked-different groups suggests that other factors were also at work. The control group showed an improvement in percent correct over the course of training, increasing from about 40% on the first block to 60% on the final block. The figures for inertial tendency suggest a plausible explanation for this learning in terms of adventitious reinforcement. Following choice responses, (which were not identified in any way in the control group), subjects tended to repeat pecks to the same colour and therefore their last response on rewarded trials was more likely than not to be the same as their initial response. The birds would thus receive adventitious reinforcement of pecks to the correct colour.

A second possibility is that subjects were sensitive to the correlation between responding and overall level of reinforcement in each session and adjusted their pattern of responding so as to maximise the number of rewards they received. Molar accounts of this kind have been proposed by several theorists (for example, Baum, 1973; Herrnstein, 1970; Nevin, 1979; Shimp, 1969).

The behaviour of the marked-different group also raises the question of how learning eventually came to emerge in that group. Although the percentage of correct responses remained around 45% over the first seven blocks of training, it did then rise to 60%. One basis for improvement is suggested by the adventitious reinforcement account outlined above. Given the presence of an inertial response tendency in the pecking behaviour of the marked-different group, the same mechanism could have operated as in the control group. A second basis for marked-different group learning could have been the molar sensitivity also described in the account of control group learning.

The penultimate issue to be considered in this discussion is the nearly significant difference between inertial tendency on trials and pseudotrials in the control group. The finding that the inertial tendency was greater during pseudotrials in control subjects was rather surprising. The difference was discovered in an analysis aimed primarily at examining the effect of the marker on inertial tendency in the marked group and it was not anticipated that there would be any divergence between trial and pseudotrial tendencies of the control group. Clearly this difference cannot be attributed to markers as none occurred in the control group. The most plausible interpretation appears to be that the receipt of food at the end of a trial somehow temporarily strengthened the tendency to repeat the most recent response. This was then reflected in a stronger inertial tendency on pseudotrials. A more refined analysis of inertial tendency, comparing pseudotrials immediately following reward with pseudotrials following nonrewarded trials, would shed light on the plausibility of this account. Two questions would still remain however, one concerning the mechanism for this effect, and the other about why the effect was restricted to the control group alone.

Marking in visual and spatial discriminations

The final section of this discussion will look at the relationship between the outcome of the current experiment involving a visual discrimination and the spatial experiment conducted by Lieberman (personal communication, September 1983) using an otherwise identical design. In most respects the visual results replicated those of the spatial version. In both experiments the marked-same group showed the highest level of learning as measured by percent correct. The number of correct responses by marked-same subjects increased substantially in both visual and spatial situations, doubling over the course of delayed reward training. Incorrect responding declined in each experiment, though the marked-same subjects maintained a higher level relative to the control group in Experiment 9 than in Lieberman's experiment.

The marked-different groups in each experiment did significantly worse than the marked-same group on the basis of percentage of correct responses. The change in number of correct responses over sessions was also similar for the two marked-different groups - each showed a slight increase in rate of responding, and each finished training at a level well below that of the marked-same group but above that of the control group. The major difference between the two marked-different groups occurred in terms of number of incorrect responses. In Lieberman's spatial experiment the rate of incorrect responding rose slightly and was significantly greater than that of the other two groups. In Experiment 9 the marked-different group showed a decline in responding to the incorrect colour which was similar to that of the marked-same group, (see Figure 20). Although towards the end of training the incorrect response rate for the marked-different

Marking in visual and spatial discriminations

The final section of this discussion will look at the relationship between the outcome of the current experiment involving a visual discrimination and the spatial experiment conducted by Lieberman (personal communication, September 1983) using an otherwise identical design. In most respects the visual results replicated those of the spatial version. In both experiments the marked-same group showed the highest level of learning as measured by percent correct. The number of correct responses by marked-same subjects increased substantially in both visual and spatial situations, doubling over the course of delayed reward training. Incorrect responding declined in each experiment, though the marked-same subjects maintained a higher level relative to the control group in Experiment 9 than in Lieberman's experiment.

The marked-different groups in each experiment did significantly worse than the marked-same group on the basis of percentage of correct responses. The change in number of correct responses over sessions was also similar for the two marked-different groups - each showed a slight increase in rate of responding, and each finished training at a level well below that of the marked-same group but above that of the control group. The major difference between the two marked-different groups occurred in terms of number of incorrect responses. In Lieberman's spatial experiment the rate of incorrect responding rose slightly and was significantly greater than that of the other two groups. In Experiment 9 the marked-different group showed a decline in responding to the incorrect colour which was similar to that of the marked-same group, (see Figure 20). Although towards the end of training the incorrect response rate for the marked-different group was greater than that of the control group, the difference was not significant in the visual discrimination experiment.

The basic relationship of the control groups to their respective marked groups was consistent across experiments. Each control group showed a reduction in the number of both correct and incorrect responses, though the drop in correct responding was more substantial in the visual experiment. On both measures of behaviour the control groups declined to a lower level than their respective marked groups. The most salient difference between the control groups was in the number of subjects that dropped out. In Lieberman's experiment, three out of fifteen control subjects stopped responding altogether and a further three reached a very low rate by the end of training. In the present experiment none of the control subjects dropped out. (The number of subjects that stopped or nearly stopped pecking in the marked groups in the two experiments was the same, and in all cases only one or two per group.) Because of the high drop out rate and variability in the control group compared with the marked groups in his experiment, (and the consequent overestimation of the controls' performance in percentage terms), Lieberman restricted his analysis to the numbers of correct and incorrect responses. The differential drop-out rate for the two control groups makes it impossible to compare them profitably on the basis of percent correct. A reason for the difference between control group drop-out rates is suggested by considering the different implications of the spatial and visual tasks for occasional, chance reinforcement. In the spatial experiment subjects were reinforced in the training stage for choice responses to the side that they had pecked least during concurrent VI training. Therefore they were more likely to respond to the incorrect side during the early sessions of the experimental phase and consequently receive few reinforcements. On the other hand, subjects in the visual experiment were rewarded to choice responses to a particular colour and position was of course irrelevant to the problem. Nonetheless subjects, like those in Lieberman's experiment, possessed side preferences. (The average percentage of responses to the preferred side on the last session of concurrent VI training was, for example, 61.) Given that the colour stimuli were alternated across sides of the key in the visual discrimination study, subjects responding on the basis of a side preference will have received occasional pairings of responses with food. Chance reinforcement of this kind may have maintained pecking behaviour by control subjects in the visual experiment in a way that was not possible in Lieberman's spatial experiment. However, if this analysis is correct, then one might expect that the number of reinforcements received in the early stages of the visual experiment would be greater than in the spatial experiment. In fact the average number of rewards received over the first five sessions by the spatial subjects, 76, was very slightly higher than the number obtained by visual subjects, 71.

Conclusion

The results of this experiment were very encouraging for the marking hypothesis. They confirmed that the failure to obtain a marking effect in earlier experiments appeared to lie with the procedures which had been adopted rather than with the visual nature of the discrimination problem. The essential similarity of the present results with those found in a spatial version of the same design (Lieberman, personal communication, September 1983) was also encouraging.

4.3 EXPERIMENT 10

The marking hypothesis was supported by the outcome of Experiment 9. However, there was also learning in the control and marked-different groups. Analysis of the data revealed an inertial tendency in the response pattern of the pigeons and this, through adventitious reinforcement, might have been the basis for learning in the control and marked-different groups.

In this experiment the intention was to examine the role of adventitious reinforcement in promoting learning and maintaining the behaviour of the subjects. This was done by introducing a differential-reinforcement-of-other-behaviour (DRO)[1] interval at the end of each trial and before the delivery of reinforcement, thereby minimizing the possibility of adventitious reinforcement. Of special interest was the impact of this upon the marked-different and control groups. A drop in the performance of these groups would provide evidence of the extent to which either (or both) were dependent upon adventitious reinforcement for the learning that occurred in Experiment 9. An additional benefit of introducing a DRO interval was that it would allow a clearer assessment of the influence of marking itself on learning. A more rigorous test of the role of marking would be made possible without the confounding effects of adventitious reinforcement.

Experiment 10 was run using the subjects from the last experiment as a preliminary test of the consequences of imposing a DRO interval. Each subject remained in the same group as in Experiment 9 and

---

[1] In a differential-reinforcement-of-other behaviour procedure reinforcement is made contingent upon the non-occurrence of a particular response. In most DRO schedules a reinforcer is delivered after a specified period has passed since the last response.

continued to be rewarded for choice responses to the same colour. It was recognised that the extensive training which the birds had already received might have produced a level of learning that was not susceptible to change as a result of the introduction of a DRO interval. Alternatively, the high response rate which had been established in the marked-same group might make it the most likely to be affected by the addition of the interval. The control group, whose responding had declined to a low level by the end of Experiment 9, might not come into contact with the DRO schedule. Despite these potential problems, the experiment was still seen as one that might provide a preliminary indication of the extent to which adventitious reinforcement was responsible for learning.

The selection of the DRO schedule was based upon evidence on changeover delays (CODs).[1] Studies by Catania (1966), Catania and Cutts (1963), and Herrnstein (1961), and a review of the literature by de Villiers (1977) suggested that a COD interval of between 1 and 3 sec is generally effective in ensuring that pigeons discriminate between schedules in a concurrent procedure. For example, in Catania and Cutts (1963) pigeons were trained in a two-key chamber with pecking initially maintained by concurrent VI schedules of reinforcement. If a pigeon was then transferred to concurrent VI-extinction schedules, pecking continued to the extinction key, albeit at a much reduced rate. The maintenance of pecking to the extinction key was attributed to accidental correlations of responses to that key with reinforcement programmed for responses to the VI key, an effect which Catania and Cutts called "concurrent supersition". However, if a COD of 1 sec

---

[1] A changeover delay specifies the minimum period that must elapse in a concurrent schedule situation before a response in one schedule is effective after a previous response in the other schedule.

was introduced, ensuring that a peck to the VI key was not rewarded until at least 1 sec had passed since the pigeon moved to that key from the extinction one, then very few responses were made to the extinction key. Applying evidence of this kind to the present situation suggested that a DRO of 1 sec would be appropriate initially to substantially reduce the benefits of adventitious reinforcement.

Method

Subjects and apparatus

The same subjects were used as in Experiment 9 except for the two birds which had dropped out during delayed reward training. The subjects were housed and maintained under the same conditions as previously, and the experiment was run using the same apparatus as in Experiment 9.

Procedure

Delayed reward training. The subjects remained in the same groups as in Experiment 9. Due to the dropping of two subjects the size of both the marked-same and marked-different groups was reduced to seven. The control group still contained eight birds. Each subject continued to be rewarded for choice responses to the same colour as in the last experiment, and the only change in training procedure was the introduction of the DRO interval. On trials where a correct choice had been made reinforcement was delivered after 10 sec only if no response had occurred during the final second of the delay interval. If a response had occurred, then the delay was extended, and continued to be extended, until 1 sec had elapsed without any response. After the fifth session of training the DRO interval was extended from 1 to 2 sec and remained at 2 sec until the end of the experiment. The DRO interval was increased because the initial interval appeared to be having little impact upon the subjects' behaviour. All subjects received 15 daily sessions which commenced on the third day after the end of Experiment 9.

Results

Percent correct

The percentage of correct responses averaged over blocks of three sessions for each of the groups is shown in Figure 23. (This and subsequent measures were based on choice and intertrial interval responding.) The figure also shows the averages for the final block of Experiment 9. There was little change in percentage correct between the first and last blocks of the experiment for any of the groups. The performance of the marked-same group was superior to that of the other groups with the marked-different group doing least well. A two-way analysis of variance confirmed that the percentage of correct responses did not change significantly during delayed reward training, $F(4, 76) = 0.16$, $p < 1$, and revealed that the Blocks × Groups interaction was not significant, $F(8, 76) = 1.07$, $p > .35$. The difference between the groups did not quite reach significance, $F(2, 19) = 3.06$, $p = 0.07$.

Number of correct responses

The data for numbers of correct responses for each group over sessions are presented in Figure 24. The figure also shows the averages for the final block in Experiment 9. The level of correct responses in the marked-same group was clearly higher than the other

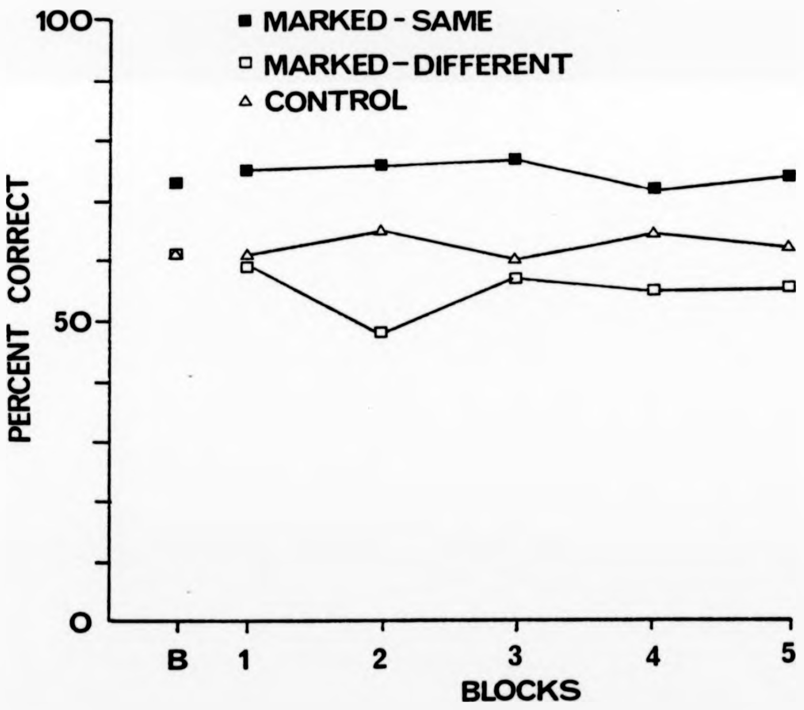

Figure 23

Mean percentage of correct responses for the final block of Experiment 9 and over blocks of three sessions for each group in Experiment 10.

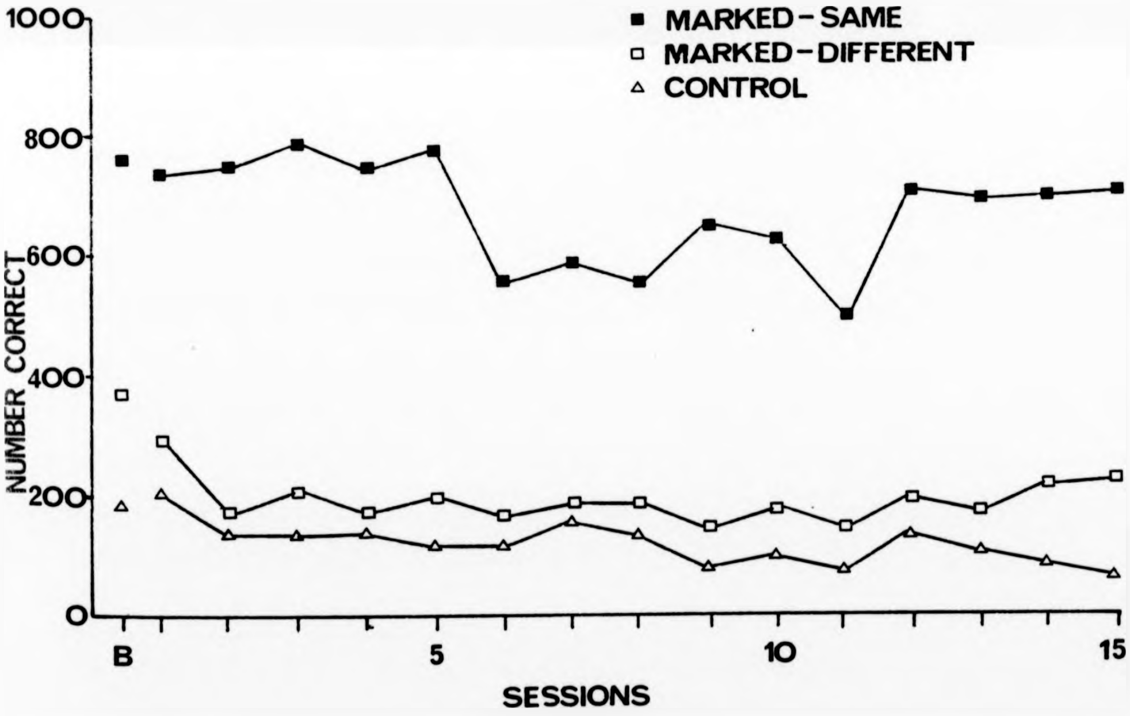

Figure 24

Mean number of correct responses for the final block of Experiment 9 and over sessions for each group in Experiment 10.

two groups which differed only slightly. A two-way analysis of variance verified that there was a significant effect of groups, $F(2, 19) = 18.44$, $p < .001$. Subsequent Newman-Keuls comparisons confirmed that the marked-same group made significantly more responses than the other groups, $p < .01$, and that the marginally higher responding of the marked-different group compared with the control subjects was not significant. Analysis of variance also revealed a significant effect of sessions, $F(14, 226) = 1.94$, $p < .05$. Inspection of Figure 24 suggested that this may reflect the drop in marked-same group responding during the middle of the experiment. Support for this impression was derived from a trend analysis of the data which showed that the quadratic trend approached significance, $F(1, 19) = 3.68$, $p = .07$. Finally, the Sessions × Groups interaction did not reach significance, $F(28, 266) = 0.79$, $p > .75$.

Number of incorrect responses

The number of incorrect responses for the three groups in each session is shown in Figure 25. Also shown in the figure are the averages for the final block of Experiment 9. The response rate in both of the marked groups was very similar and showed an initial decline followed by a recovery towards the end of training. The control group responded at a lower level and did not increase its rate during the final blocks. Analysis of variance revealed the number of incorrect responses was not in fact significantly higher in the marked groups than in the control group, $F(2, 19) = 2.65$, $p < .10$. The number of responses did vary over sessions, $F(14, 266) = 2.61$, $p < .01$ and once again this appeared to arise from the low point in responding about two-thirds of the way through the experiment. Trend analysis supported this interpretation, as the quadratic trend was significant, $F(1, 19) = 9.18$, $p < .01$.

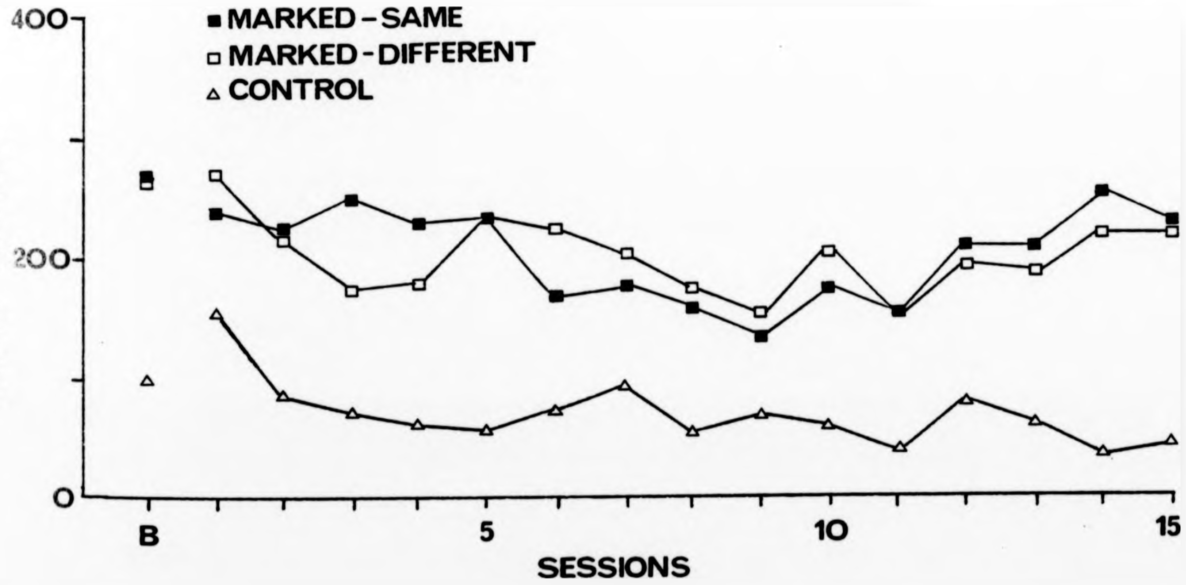

Figure 25

Mean number of incorrect responses for the final block of Experiment 9 and over sessions for each group in Experiment 10.

Discussion

The aim of this experiment was to assess the role of adventitious reinforcement in learning, particularly in the control and marked-different groups. The introduction of a 1-sec DRO did not appear to make very much of an impact upon the subjects' behaviour during the first few sessions and it was accordingly increased to 2 sec after the fifth session.

It was anticipated that the control group would suffer most from the elimination of adventitious reinforcement and to some extent this prediction was borneout by the results. The average number of correct responses by control subjects on the final block of Experiment 9 was 190 and by the end of the present experiment this had declined to 70. The number of incorrect responses also dropped between the final sessions of Experiments 9 and 10, from 100 to 47. However, the percentage of correct responses remained the same at 61%. On the positive side though, it should be noted that after the addition of the DRO contingency the general trend of an increase in percent correct by the control group which was apparent in Experiment 9, (see Figure 18), halted. The level of learning failed to reach that shown by the marked-same group.

The marked-different group was also expected to be affected by the DRO interval. The number of correct responses showed a sharp drop between the last block of Experiment 9 and the first session of Experiment 10 - responses went from 370 to 300. This decline continued during part of the current experiment but then recovered slightly to reach about 230 responses by the final session. A drop in the number of incorrect responses also occurred over the first half of this experiment and was followed by an increase towards the end. In the final session the response rate was 220 compared with a figure of 280 in the last block of Experiment 9. In the marked-different group the percentage of correct responses fell as predicted though not by a substantial amount. Comparing the final block of the last experiment with the final session of this experiment shows a fall of five percentage points.

As expected the prevention of adventitious reinforcement had least impact on the performance of the marked-same group. By the end of training the number of correct responses, 710, was similar to the level of 760 at the end of Experiment 9, though there was a fairly large drop halfway through training. The number of incorrect responses fell more significantly from 280 to 230 at the conclusions of Experiments 9 and 10 respectively. There was no decline in percentage of correct responses which remained around 75%, (though it should be noted that the upwards trend in Experiment 9 did not continue).

Precise interpretation of the significance of this pattern of results is not easy. The amount of training that subjects had already received may have meant that the impact of eliminating adventitious reinforcement was not the same as it would have been had the DRO contingency been introduced with naive subjects. It is difficult to be sure whether particular trends that were apparent in Experiment 9, (for instance, the increase in percent correct shown by the marked-same group and illustrated in Figure 18), would have continued had the DRO interval not been added. On the other hand, it is possible that the decline, for instance, in both correct and incorrect responding by the control group in Experiment 10 was simply a continuation of trends already obvious in the last experiment and shown in Figures 19 and 20.

Another problem is the recovery in the level of both correct and incorrect responding in the marked groups after an initial fall. One possibility is that the elimination of adventitious reward had an initial impact that was eventually offset by a continued general effect of marking upon pecking.

On balance though, despite these difficulties of interpretation, the broad picture presented by the outcome of Experiment 10 is reasonably consistent with the suggestion that adventitious reinforcement played a part in the learning seen in the last experiment, especially for the control and marked-different groups. Thus the results provide some support for the adventitious reinforcement interpretation.

Another problem is the recovery in the level of both correct and incorrect responding in the marked groups after an initial fall. One possibility is that the elimination of adventitious reward had an initial impact that was eventually offset by a continued general effect of marking upon pecking.

On balance though, despite these difficulties of interpretation, the broad picture presented by the outcome of Experiment 10 is reasonably consistent with the suggestion that adventitious reinforcement played a part in the learning seen in the last experiment, especially for the control and marked-different groups. Thus the results provide some support for the adventitious reinforcement interpretation.

4.4 EXPERIMENT 11

Experiment 9 provided evidence that marking is a phenomenon that occurs in visual as well as spatial situations. Analysis of the data also showed that an inertial response tendency could have produced learning through chance pairings of correct responses with reinforcement. This adventitious reinforcement might have obscured the extent of the effect of marking, in particular by promoting a higher level of learning in the control and marked-different groups than would otherwise have occurred.

In Experiment 10 a DRO interval was introduced in order to remove adventitious reinforcement. The effects of this contingency were difficult to predict and interpret as the subjects had already had substantial experience with the discrimination. Therefore in this experiment the same design was used except that the DRO contingency was introduced from the beginning of delayed reward training.

Once again three groups of subjects were trained - marked-same, marked-different, and control. The aim was to test the hypothesis that adventitious reinforcement provided the basis for learning in the control group and, possibly to a lesser extent, in the marked-different group. Although it was expected that the DRO interval might have some impact on response level and learning in the marked-same group, the marking hypothesis would still predict that the group's performance should be better than that of the other groups and that, with adventitious reinforcement substantially reduced, a marking effect should be more apparent.

Method

Subjects

The subjects were 27 adult Rock pigeons which were experimentally naive at the start of the experiment. They were housed and maintained under the same conditions as in previous experiments.

Apparatus

The same apparatus was used as in Experiment 9.

Procedure

Magazine training, shaping and pretraining. The same procedure was used as in Experiment 9. Four subjects received two sessions of concurrent VI 40-sec and only one session of concurrent VI 60-sec.

Delayed reward training. The subjects were divided into three groups of nine subjects. As in Experiments 9 and 10, the groups were called marked-same, marked-different and control.

The training procedure was the same as in Experiment 9 except for two changes. One was the introduction of a 2-sec DRO interval to prevent adventitious reinforcement. The first response of each trial was designated the choice response. Reinforcement was scheduled for delivery 10 sec after a correct choice though it was delayed until 2 sec had elapsed since the last peck. The second change was that subjects received 30, rather than 20, daily sessions of delayed reward training.

Results

Percent of correct responses

The percentage of correct responses averaged over blocks of four sessions for each group is shown in Figure 26. However analysis of variance was not performed on these data for two reasons. The first was that on slightly over 10% of individual subject/sessions, the subject made fewer than the 10 responses which had been set as a minimum criterion for the calculation of percent correct. The second was that the distribution of these very low response subject/sessions across groups was uneven with over twice as many occurring in the control group as in the marked groups put together. Therefore the percentage of correct responses was not a reliable measure of performance in this experiment and the analysis focused on the number of correct and incorrect responses.

Number of correct responses

The numbers of correct responses averaged over three-session blocks for each of the three groups are illustrated in Figure 27. All groups showed a sharp decline in responding during the first three blocks. This fall continued in the control group but was reversed in the two marked groups. By the end of training the level of responding in the marked-same group was approximately twice that in the marked-different group. A two-way analysis of variance was carried out on the data presented in Figure 27 for the factors of experimental treatments and trials. It showed that the groups differed significantly in the number of correct responses, $F(2, 24) = 7.36$, $p < .01$, and that the level of correct responding varied across blocks, $F(9, 216) = 8.00$, $p < .0001$. The Blocks × Groups interaction was also significant, $F(18, 216) = 5.00$, $p < .0001$, confirming that the groups differed over blocks. A simple main

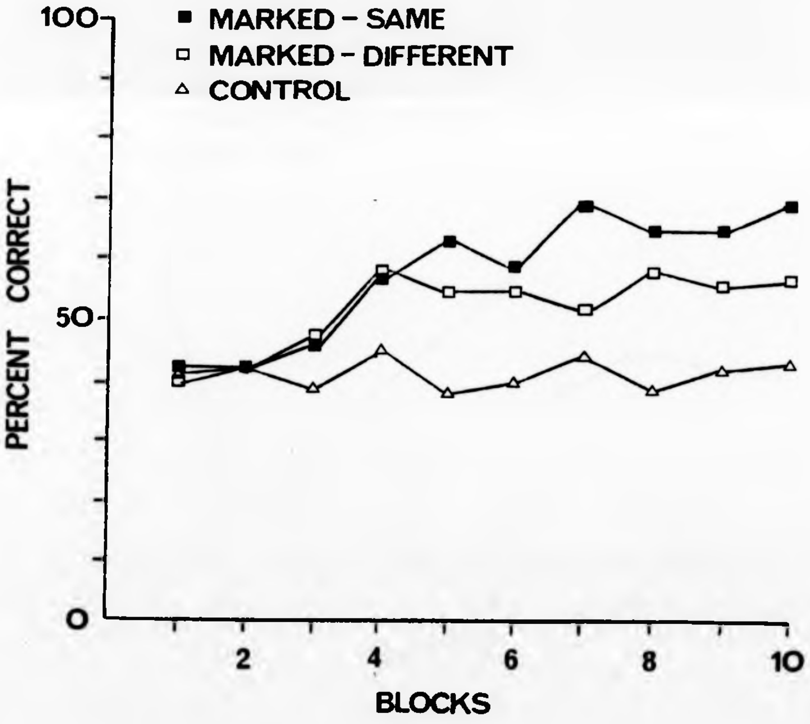

Figure 26

Mean percentage of correct responses over blocks of three sessions for each group in Experiment 11.

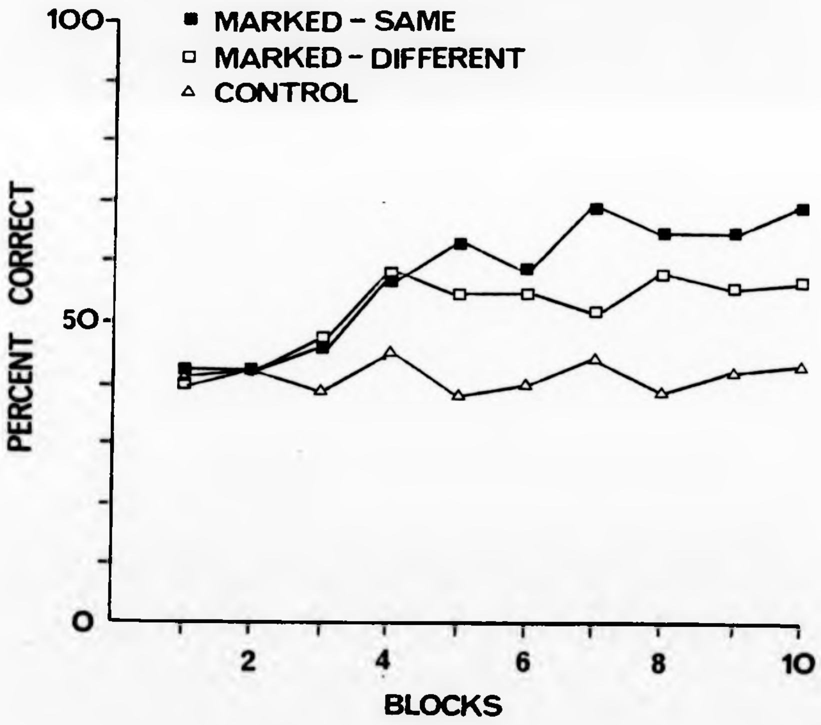

Figure 26

Mean percentage of correct responses over blocks of three sessions for each group in Experiment 11.

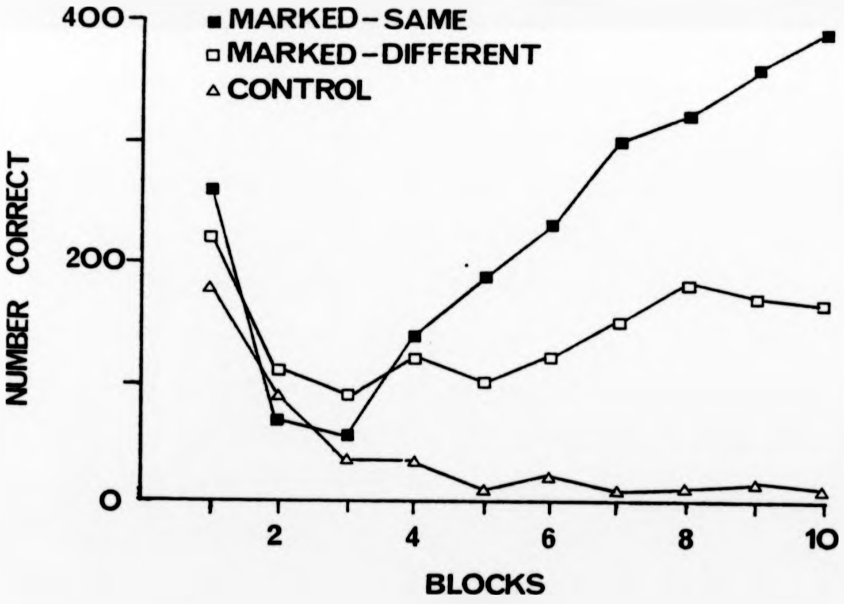

Figure 27

Mean number of correct responses over blocks of three sessions for each group in Experiment 11.

effects analysis showed that the groups differed on Blocks 6 to 10 inclusively, *F*s(2, 24) > 7.24, *p*s < .01, but not on the remaining blocks. Subsequent Newman-Keuls comparisons indicated that on Block 6 responding was significantly higher in the marked-same group than in both the marked-different group, *p* < .05 and the control group *p* < .01. On Blocks 7 to 10 the marked-same group responded more than either of the other groups, *p* < .01, and the marked-different group had a higher response rate than the control group, *p* < .01.

Number of incorrect responses

Data for the number of incorrect responses averaged over blocks of three sessions for each group are presented in Figure 28. Once again there was a substantial fall in the number of responses by all groups over the first three blocks. In subsequent blocks the downward trend in responding was maintained only in the control group. Both marked groups, which did not differ from one another, showed a levelling out of responding and then an increase over the second half of training. The data were subjected to a two-way analysis of variance which revealed that the groups differed significantly in number of incorrect responses, *F*(2, 24) = 4.42, *p* < .05, and that the level of responding changed across blocks, *F*(9, 216) = 25.83, *p* < .0001. The interaction between blocks and groups was not significant, *F*(18, 216) = 1.33, *p* > .15. Newman-Keuls comparisons were subsequently made of the differences between groups and these showed that incorrect responding was significantly lower in the control group than either the marked-different group, *p* < .01, or the marked-same group, *p* < .05.

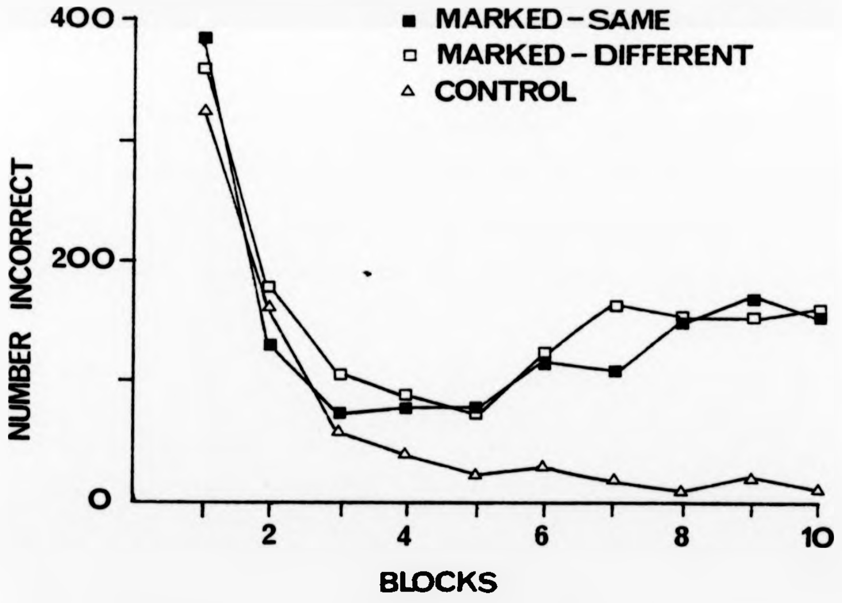

Figure 28

Mean number of incorrect responses over blocks of three sessions for each group in Experiment 11.

Discussion

The objective of this experiment was to examine the role that adventitious reinforcement had played in promoting learning in Experiment 9. The introduction of a DRO contingency rendered adventitious reinforcement unlikely and allowed a more direct assessment of the effect of marking.

A comparison of the total number of responses in Experiments 9 and 11 shows a lower level of responding in all groups in the second of these experiments. This difference, which is apparent both for correct responses (see Figures 19 and 27) and for incorrect responses (see Figures 20 and 28), is found throughout the course of training. Several factors could be responsible for the lower number of responses in Experiment 11.

One possibility is that the baseline response rate in the two experiments differed. Analyses of the data for the final session of concurrent VI training show that this is not a satisfactory explanation as the average rate in the present experiment was in fact marginally higher than in Experiment 9, the figures being 1500 and 1300 respectively.

A second alternative is that the time period over which responses were being recorded in this experiment was shorter. This could have occurred if the trial duration had increased somewhat as a result of the addition of the DRO contingency, thereby reducing the intertrial interval period on which responses were based. However, evidence from the total intertrial interval time on Session 1 in each of the experiments provides little support for this view. (Session 1 was selected as the response rate was highest on that session in Experiment 11 and any resultant decrease in the intertrial interval time most substantial.) In Experiment 9 the total intertrial interval times for the marked-same, marked-different and control groups were 1500, 1500 and 1490 sec compared with 1401, 1412 and 1384 sec respectively in Experment 11. Thus even on the session when the intertrial interval procedure over which responses were recorded was reduced most significantly relative to Experiment 9, the reduction only amounted to about 7%. Therefore the change in the length of the period on which the response measures are based is far too slight to account for the observed fall in number of responses.

The third and most plausible explanation of the lower number of responses in the current experiment is the impact of the DRO interval and the resultant elimination of adventitious reinforcement. No chance pairings of responses and food occurred which could help maintain responding at a higher level (Dews, 1960).

A fourth factor that may have been operating as well as the DRO contingency was an increase in the delay of reinforcement itself. If responses did occur in the final two seconds of a trial and thereby bring the DRO contingency into operation, then the resultant increase in trial duration may have had an effect over and above that of the DRO interval. Evidence that unsignalled delays of reinforcement can lower the response rates of pigeons in free-operant situations has been provided by several investigators (Richards, 1981; Sizemore and Lattal, 1978; Williams, 1976). (Unsignalled delay of reinforcement studies are more directly comparable to the present design than signalled ones which normally involve a complete darkening of the chamber for the duration of the delay.) However, these studies also suggest that a floor effect is found with delays over the range 4 to 15 sec and possibly longer. The average trial durations for each group on Session 1 were 12.1, 12.5 and 13.0 sec for the marked-same, marked-different and control groups respectively. Given that the DRO contingency in this experiment did not extend the delay more than three seconds at most beyond the scheduled 10 sec, it seems unlikely that the reduction in response rates can be attributed to an increase in actual delay of reinforcement.

When the results of Experiment 11 are compared with those of Experiment 9 they suggest that a substantial reduction in adventitious reinforcement was responsible for a general lowering of the response rate. The performance of individual groups can now be compared in order to assess the contributions of adventitious reinforcement and marking. In both experiments the level of learning was highest in the marked-same group. In this experiment, as in the previous one, the number of correct responses by marked-same subjects was significantly greater than in the other groups (see Figures 19 and 27). It is apparent that chance pairings of food with correct responses did make a contribution to marked-same group learning in Experiment 9. The substantial fall in number of pecks during the first three blocks in the present experiment was not seen in Experiment 9, and the level of responding reached as early as Session 12 in Experiment 9 was not reached even after thirty sessions in Experiment 11. Indeed the rate at the end of training in the latter experiment was only half that in the former. However, despite this effect of adventitious reinforcement, it is clear that marking was also at work. Although the number of incorrect responses in the marked-same group was as high as in the marked-different group, the substantially higher number of correct responses in the marked-same subjects in Experiment 11 shows that learning was facilitated by following responses by a salient event. The increase in the number of incorrect responses after an initial drop is a problem that will be returned to in the discussion of the marked-different group.

The introduction of a DRO interval had a strong effect on the performance of the control group. As predicted, the elimination of adventitious reinforcement substantially reduced both the number of responses and the level of learning (as measured by the difference between correct and incorrect response rates). In Experiment 9 the average numbers of correct and incorrect responses in the final block were 190 and 100 respectively. The equivalent figures for Experiment 11 were 10 and 10. This suggests that the primary basis for control-group learning in Experiment 9 had been adventitious reinforcement.

The problem of the level of incorrect responding in the marked-same group has already been identified - another aspect of the outcome of this experiment that requires careful explanation is the performance of the marked-different group. Application of the marking hypothesis would suggest that the level of correct responding should have declined in marked-different subjects because incorrect responses were marked on reinforced trials. Two predictions could therefore be derived from the marking hypothesis. First, the level of correct responding should be lower than the level of incorrect responding in the marked-different group. In neither Experiment 9 or 11 was this prediction fulfilled - in both the levels of correct and incorrect responding were about the same. A second, though weaker, prediction is that the number of correct responses in the marked-different and control groups might be similar. In Experiment 9, the number of correct responses in the marked-different group was in fact higher than in the control group towards the end of training (see Figure 19). However, this difference was not significant and it was therefore decided to defer discussion of the possible reasons for it. In Experiment 11, the number of correct responses was once again greater in the marked-different group, and this time the difference was significant. This replication of a substantial rate of correct responses in the marked-different group, (both relative to incorrect responding in that group and to correct responding by control groups), suggests that it requires explanation, particularly as it occurred in an experiment free from the confounding factor of adventitious reinforcement.

One explanation for the effect may lie with the invisible-trial procedure used in this experiment. In the earlier, discrete-trial and maze experiments (Fuller, 1981; Lieberman, personal communication, 1981; Lieberman et al., 1979; Thomas et al., 1983) the absence of reinforcement at the end of a trial following an incorrect choice was an easily discriminable event. As such, non-reinforcement could be expected to have initiated a memory search that would identify the marked choice response and facilitate learning about it. However, with the invisible-trial design learning about choices on non-rewarded trials is much more difficult as incorrect choices have no readily identifiable consequences at the end of the delay. (On correct trials the two procedures are much more similar as in both cases the presentation of food provides a significant event of the kind likely to start a backward scan through memory store.) Thus in the marked-different group being considered here, the correct response which was marked on non-rewarded trials would be less likely to be associated with non-reinforcement. Correct responding might therefore have been maintained at a higher rate than would have been the case in an effective discrete-trial procedure.

A second explanation for the maintenance of correct responding in the marked-different group is suggested by consideration of the role of the marker. So far the marker's function on rewarded trials has been discussed in terms of whether a correct, or incorrect choice (depending upon the group) would be identified by a food-initiated memory search. However, additional ways in which the marker might have operated can be identified that would have kept up the level of correct responding in the marked-different group. One possibility is that an effect of the marker was simply to draw attention to the act of pecking itself rather than as a response to a particular colour. The backward scan through memory following food may have promoted learning by marked-different subjects about pecking behaviour in general and reward, instead of (or as well as) about incorrect responses and reward. This process could be called response marking in contrast to the choice marking discussed up until now.

A second possibility suggested by a broader view of the role of marking is based on evidence that pigeons can learn about sequences of responses as well as single responses (see, for example, Fetterman and Stubbs, 1982; Grayson and Wasserman, 1979). If marking of response sequences could occur, then a pattern of alternating responses would be strengthened in the marked-different group, thereby helping to maintain responding to the correct colour. Alternation would be strengthened because the sequence of pecks which preceded a marker necessarily involved a change from correct to incorrect responding or vice versa. This hypothesis could be tested by comparing the level of intertial tendency in the two marked groups - one would expect a greater inertial tendency in the marked-same group as in that group sequences involving continued pecks to the same colour would be marked. However, analysis of variance for the Experiment 9 data showed that in that experiment the two groups did not differ in inertial tendency, a result which does not support the response-sequence marking hypothesis.

A third explanation is available for the level of correct responding in the marked-different group relative to the control group. It is based on the concept of "sensory reinforcement" used by Kish (1966, p.110) to refer to "a primary reinforcement process resulting from the response-contingent presentation or removal of stimuli of moderate intensity" which cannot be subsumed under the categories of primary, or secondary, appetitive and aversive reinforcement. Kish (1966) reviewed a number of studies of sensory reinforcement which showed that it could produce behavioural changes if made contingent upon responses. For example, he reported that rats sometimes learned to lever-press when the only result of the response was to turn on a brief light. It may be that the marker used in this experiment also had sensory reinforcing properties. If this was the case, then a sensory reinforcement account might help explain why correct responding (and, of course, incorrect responding which was also marked) was higher in the marked-different group than the control group.

This discussion has therefore identified three processes that might have been responsible for the fact that the number of correct responses in the marked-different group was double that in the control group toward the end of training. One possibility is that marking of responding generally can occur as well as the marking of choice responses to a particular colour. A second possibility is that marking failed to reduce correct responding because the effects of marking occur primarily on rewarded trials when an invisible-trials procedure is used. A third hypothesis is that the marker acted as a sensory reinforcer and thereby helped maintain correct (and incorrect) responding at a level above that in the control group.

As was noted earlier in this discussion, the number of incorrect responses in the marked-same group was the same as in the marked-different group and significantly higher than in the control group (see Figure 27). According to the marking hypothesis one might expect that incorrect responding by marked-same birds should fall below that in marked-different subjects and reach a level close to that of the control group. The reasons why this does not happen are suggested by the analysis above of correct responding in the marked-different group. One is a general facilitation of responding by marking, another is the possibility that the marking effect occurs primarily on rewarded trials when an invisible-trials procedure is adopted, and a third is that the marker provides sensory reinforcement. All three factors could have kept up incorrect responding by marked-same subjects. A consequence of this would, incidentally, have been to depress the performance of the marked-same group in Experiment 9 as measured by percent correct and bring it to a level closer to that found in the control group.[1]

This discussion has covered the general effect of the elimination of adventitious reinforcement, the implications of the results for marking, and some additional explanations of certain aspects of the results. The remainder of the discussion will involve a comparison of the outcome of this experiment with that of a subsequent spatial replication by D. A. Lieberman (personal communication, January 1984) The general pattern of results in Lieberman's experiment was similar to that found in Experiment 11 though a comparison of the two studies implies that marking may be stronger in the spatial situation. This

---

[1] See footnote 2 on page 178.

is suggested by two aspects of the results. The first is that the number of correct responses by the marked-same group reached a higher level in Lieberman's experiment than in the current one. By the final block of the spatial experiment the response rate was about three times its level in the first block, whereas in the visual experiment the rate only doubled. This difference in performance cannot be attributed to differences in either the initial response rate (which was similar in both studies) or the length of training (which was in fact shorter in the spatial experiment). However, it should also be noted that the level of correct responding by control subjects was also higher in Lieberman's study than in the present experiment. The average rate for the final three sessions of Lieberman's experiment was 70 responses compared with only 10 responses after a similar number of sessions in this experiment. It is therefore possible that this apparent divergence in the strength of the marking effect reflects some more basic difference between responding in spatial and visual situations. A final alternative is that the difference in results is a chance effect attributable to random variation.

A second feature of the results that suggests a stronger spatial than visual marking effect is the number of incorrect responses in the marked-different group relative to the marked-same group. In Experiment 11, the number of incorrect responses in the two marked groups did not differ, in contrast with Lieberman's experiment where incorrect responding in the marked-different group was significantly higher than in the marked-same group. (Inspection of the data shows that the level of incorrect responding in the marked-same group was about the same in Lieberman's experiment and Experiment 11. The divergence between the two experiments in terms of the relative performance of the two marked groups' level of incorrect responding, is therefore attributable to a higher rate in Lieberman's marked-different group.) This finding implies that the effect of marking an incorrect response on rewarded trials is more substantial in spatial discriminations.

---

[1] The almost identical number of incorrect responses on the final block in the two marked groups makes it unlikely that the difference in correct responses is due to the marker producing a general increase in responding in the marked-same group.

[2] A fourth potential explanation of the higher level of responding in the marked groups relative to the control group is that the marker became associated with reward and acted as a secondary reinforcer for the response of pecking.

4.5 GENERAL SUMMARY AND DISCUSSION

The major aim of the visual discrimination experiments reported in this chapter was to test whether the failure to find marking in Experiments 1 to 8 reflected a limit to the generality of the phenomenon, or arose from the procedures which had been used. The successful demonstration of marking in Experiments 9 to 11 with an invisible-trials procedure and the same red-green discrimination as in earlier studies shows that the primary reason for previous failures lies with the discrete-trial design. In all of these experiments the number of correct responses by the marked-same group was significantly greater than in the other groups.

Arousal and generalisation decrement accounts of the superior performance of the marked-same subjects can be rejected. Both hypotheses would predict that the level of learning found in the marked-same and marked-different groups should be similar. In none of Experiments 9, 10 or 11 was this the case and the results of these studies can therefore be added to past demonstrations of the inadequacy of arousal accounts of the marker (Lieberman, personal communication, September 1983; Lieberman et al., 1979; Thomas et al., 1983).

The presence of an inertial response tendency was detected in Experiment 9 and its role in learning was tested in the subsequent experiments. Experiment 10, which involved a substantial reduction in adventitious reinforcement by the introduction of a DRO interval, provided preliminary evidence that chance pairings of correct responses with food had played a part in learning in Experiment 9. A second study with a DRO contingency, Experiment 11, showed that adventitious reinforcement had been responsible for virtually all the learning shown by control subjects in the first visual invisible-trials experiment. It also demonstrated that adventitious learning had made some contribution to the performance of the marked groups.

The general pattern of results in these experiments, and in particular the number of incorrect responses by marked-same subjects and correct responses by marked-different subjects, suggests two other aspects to the marking process. The first is that the benefits of marking arise primarily on rewarded trials when an invisible-trials procedure is adopted. When reinforcement is received it initiates a scan of the subject's memory store and this search is likely to identify the marked response as a causal or predictive factor in relation to food. With an invisible-trials procedure non-reinforcement is less likely to produce a memory search. This situation contrasts with that in a discrete-trials procedure where the non-delivery of food at the end of a trial is probably much more obvious. This analysis implies that subjects are not likely to learn about marked responses on trials where an incorrect choice occurs, and may therefore account in part for the maintenance of incorrect responding in the marked-same group and correct responding in the marked-different group.

A second possibility raised by the results of these experiments is that the effect of marking is not only to facilitate learning about correct and incorrect choices, but also about responding generally. In some instances a backward scan of memory may fail to identify that a peck occurred to, say, green, but will nonetheless identify that a peck as such occurred. Thus the generally higher level of responding in the marked groups may reflect a general effect of marking upon pecking.

One other way in which the marker may have operated in these experiments is as sensory reinforcer (Kish, 1966). As such, the effect of the marker might have been to contribute to the higher level of responding found in these experiments in the marked groups relative to the control groups.

Some contribution to learning, particularly in the marked-different group, may have arisen from a molar sensitivity by subjects to the relation between their pattern of responding and the overall level of reinforcement received in a session (Baum, 1973; Herrnstein, 1970; Nevin, 1979; Shimp, 1969).

A final point to emerge from the series of experiments reported in this chapter concerns the comparison between visual and spatial marking. It is possible to evaluate the effects of marking in visual and spatial discriminations because the designs of Experiments 9 and 11, and Lieberman (personal communication, September 1983) and Lieberman (personal communication, January 1984) respectively were very similar. It is clear from these studies that marking produces essentially the same effects in both types of discrimination. However, the effect does appear to be stronger in spatial tasks. In Experiments 9 and 11, the level of incorrect responding in the marked-same and marked-different group was the same. In contrast, in Lieberman's experiments the number of incorrect responses was greater in the marked-different subjects. As incorrect responding was marked on food trials in the marked-different groups, these results imply that the effect of the marker was greater in the spatial studies.

Better spatial marking was also suggested, though to a lesser extent, by the level of correct responding in the marked-same groups. The rates found during the final blocks of Lieberman's experiments were higher than in Experiments 9 and 11. However, the initial rate at the start of training in Experiment 9 was lower than in the comparable spatial experiment (Lieberman, personal communication, September 1983) and the difference in terminal rates may reflect this rather than a more powerful spatial marking effect. The evidence provided by a comparison of Experiment 11 and Lieberman's (personal communication, January 1984) experiment is stronger. In the visual study the rate of correct responding doubled over the course of training whereas in the spatial study it trebled. The initial response rate in the two experiments was similar.

The question of why marking might be greater in spatial than in visual discriminations will be considered in detail in the final chapter of the thesis. One possibility, however, is that the difference in the marking effect reflects the relative salience or novelty, or both, of the two markers used in the experiments just discussed. In Experiments 9 and 11, the marker was a 1-sec change from red and green key illumination to white. Lieberman's experiments involved a change from white to red and green key illumination, also for 1 sec, as the marker and intuitively it appears at least plausible that this represented a more salient or novel event.

CHAPTER 5

CONCLUSION

5.1 GENERAL INTRODUCTION

The purpose of this research was to explore the phenomenon reported by Lieberman et al. (1979) and Thomas et al. (1983). Lieberman and his colleagues found an enhancement of learning in a two-choice spatial discrimination task with rats if every choice response was followed by a salient event. They proposed a marking hypothesis to account for this result. The salient event, (which was either being handled, a burst of white noise, or a flash of light), caused the subject to scan its memory for predictive or causal cues. The cue most likely to be identified by this search was the choice response that had immediately preceded it and which was therefore marked in the animal's memory by the additional attention and rehearsal it received. Later, when reward was given, a second backward scan through the memory store was triggered and the marked choice response was identified as the causal factor.

The research conducted for this thesis had three major aims. The first was to demonstrate the generality of the marking phenomenon. All the experiments carried out by Lieberman et al. (1979) and Thomas et al. (1983) involved rats trained on spatial discriminations in mazes. The experiments reported here used pigeons in operant chambers and a visual discrimination problem. The results show that marking can facilitate learning by pigeons in a visual, operant discrimination but that an effect is not always found. The second objective was to develop an automated procedure for studying marking. Such a procedure would eliminate any possibility of experimenter bias in the running of experiments. It would also be less time-consuming than work involving mazes. In this objective the research was successful. The third aim of this research was to investigate the marking effect more thoroughly once a satisfactory paradigm had been found. Unfortunately the extent of this investigation was constrained by the difficulties in obtaining an effect in many of the experiments.

The next section of this chapter will summarise the results of the experiments. This will be followed by a list of the conclusions that can be drawn from this research. The remainder of the chapter will be devoted to exploring the implications of this work and the relationship between marking and other areas of research in animal learning.

would also be less time-consuming than work involving mazes. In this objective the research was successful. The third aim of this research was to investigate the marking effect more thoroughly once a satisfactory paradigm had been found. Unfortunately the extent of this investigation was constrained by the difficulties in obtaining an effect in many of the experiments.

The next section of this chapter will summarise the results of the experiments. This will be followed by a list of the conclusions that can be drawn from this research. The remainder of the chapter will be devoted to exploring the implications of this work and the relationship between marking and other areas of research in animal learning.

5.2 SUMMARY OF RESULTS

All the experiments reported here involved the same red-green discrimination. The basic results were as follows.

Discrete-trial I experiments

In Experiments 1 and 2 a 15-sec delay of reward was used. In the first experiment choice responses were marked with a 2-sec offset of the houselight and in the second with a combination of houselight offset and a burst of white noise, also for 2 sec. In each experiment both marked and control groups learned equally well. However, in both groups the choice response turned off the illumination on the non-chosen side of the key and this might have inadvertently acted as a powerful marker, making the houselight-offset and white-noise markers redundant.

In Experiment 3 this source of non-differential marking was eliminated by not turning off either side of the key after a choice response. The marked subjects showed no learning despite choice responses being followed by houselight offset and white noise. In the control group a little learning was found. This result implied that the offset of the non-chosen side of the key might have been responsible for the learning observed in Experiments 1 and 2.

Experiment 4 examined more directly the role of the darkening of the non-chosen side of the key. Did it act as a salient event that initiated a memory search, or did it facilitate learning through a non-marking process involving perceptual focusing as a result of additional exposure to the chosen colour? Three groups were used. Learning in the whole-key-marker group - where the key was completely darkened immediately after a choice response - and in the control group did not differ. The best performance was found in the half-key-marker group. From this result it was concluded that learning had been produced in this group of experiments as a consequence of perceptual focusing rather than marking. Following choice responses with a burst of white noise or the offset of the houselight, (or keylight change as such), did not facilitate learning relative to control subjects. Habituation to the markers or elements of the markers might have been responsible for their failure to promote additional learning in Experiments 1 to 3.

Discrete-trial II experiments

One of the reasons for the failure to obtain marking in the first set of experiments might have been that the problem used was simply too hard, and so in the next series of experiments the delay of reinforcement was reduced from 15 to 6 sec. In addition, the duration of the marker was changed from 2 to 0.5 sec in order to minimise the possibility of marking more than one peck to the key.

In Experiment 5 the marked group, in which choice responses produced a darkening of both keylight and houselight for 0.5 sec, did consistently better than the control group. However, the difference in performance was not significant.

Experiment 6 was an attempt to improve upon the design used in the preceding experiment. A new marker was adopted in which key illumination changed from red and green to amber for 0.5 sec following a choice response. In all previous experiments the key illumination in both groups changed from red and green to white at some point during the delay interval. This change might have functioned as an adventitious marker, thereby depressing marked group performance and enhancing control group learning. This possibility was eliminated in Experiment 6 by maintaining red and green stimuli until the end of the delay. Learning was found in both groups but the control group did significantly better than the marked group - a counter-marking effect. Adventitious reinforcement produced by an inertial response tendency was seen as a probable basis for control group learning. The poor performance of the marked group might have been produced by the amber marker interfering with memory for the discriminative stimuli, or by the marker disrupting inertial responding.

Experiment 7 incorporated the markers from Experiment 5 - the blackout - and Experiment 6 - the amber flash. The most substantial improvement in performance was found in the blackout-marker and control groups, with blackout-marker subjects doing marginally better towards the end of training. The increase in percent correct was lowest in the amber-marker group, though this result was difficult to interpret as its starting point also differed from the other two groups. The results suggested, however, that the amber marker might be peculiarly damaging to learning because of colour interference.

Experiment 8 involved an attempt to enhance the effect found in Experiment 5. The change in key illumination during the delay was reintroduced. The same marker as in Experiment 5 was used except that a 0.5 sec flash of the magazine light was added in order to enhance the marker's salience. Learning was found in both groups and the marked group did not differ from the control group.

The results of these experiments suggested that either the procedure was inappropriate in some way, or that marking was limited to spatial tasks. The purpose of the next set of experiments was to clarify the issue by investigating visual discrimination learning using a different procedure.

Invisible-trial experiments

In these experiments trials were not identified as such by the onset and offset of the discriminative stimuli. The key remained lit with the red and green stimuli for the entire session except during markers. The marker used was a change in key illumination from red and green to white for 1 sec. Markers were presented during a window from 2 to 9 sec after a choice response. Three groups were employed in each experiment. In the marked-same groups, the marked response was the same as the choice response. In the marked-different groups, the marked response was different from the choice, and in the control groups, no markers were presented.

In Experiment 9, the marked-same group did substantially better than the other two groups. Learning was, however, also found in the control and marked-different groups. An inertial response tendency was found in all groups.

Experiment 10 was a preliminary investigation of whether control and marked-different subjects might have learned in the preceding experiment as a result of adventitious reinforcement. A DRO interval was introduced to reduce substantially chance pairings of responses with food. The results were in line with predictions derived from the adventitious reinforcement hypothesis.

Experiment 11 replicated the last experiment using naive subjects. It confirmed that adventitious reinforcement was the major factor determining control group behaviour and that it played a part in the performance of the marked groups. The marked-same group clearly showed, however, that a marking effect was also in operation.

These experiments indicated that visual marking occurs and that the primary problem in Experiments 1 to 8 was the procedure used.

Experiment 11 replicated the last experiment using naive subjects. It confirmed that adventitious reinforcement was the major factor determining control group behaviour and that it played a part in the performance of the marked groups. The marked-same group clearly showed, however, that a marking effect was also in operation.

These experiments indicated that visual marking occurs and that the primary problem in Experiments 1 to 8 was the procedure used.

5.3 SUMMARY OF CONCLUSIONS

The major conclusions which can be drawn from this research are summarised below. The experiments that provide support for each conclusion are listed in brackets.

1. Marking is found in visual discriminations (Experiments 9 to 11).
2. Marking can occur when pigeons are used as subjects (Experiments 9 to 11).
3. Marking can be demonstrated in the operant chamber (Experiments 9 to 11).
4. Following choices with a salient event appears to mark not only the particular choice response but also responding more generally (Experiments 9 to 11).
5. Marking may occur primarily on rewarded trials when an invisible-trials procedure is used (Experiments 9 to 11).
6. Visual marking may be weaker than spatial marking (Experiments 9 to 11 and Lieberman, personal communications, September 1983, January 1984).
7. Marking does not always facilitate learning (Experiments 1 to 8).
8. A counter-marking effect is sometimes found when a choice response is followed by a marker (Experiment 6, maybe Experiment 7, and possibly Experiment 3). This effect might be attributable to colour interference between the stimulus used as the marker and the discriminative stimuli.

9. Pigeons can show an inertial tendency in their responding (Experiment 9) which leads to adventitious reinforcement. This can provide a basis for discrimination learning with delayed reward (Experiments 7, 9, 10 and 11).

This research therefore demonstrates the broad generality of the marking phenomenon whilst suggesting limits to the circumstances in which an effect will be found. Several questions are raised by the results presented here, both in relation to one another, and in contrast with the findings of other investigators. These questions will be addressed in the remaining sections of this chapter.

5.4 DISCRETE-TRIAL AND INVISIBLE-TRIAL PROCEDURES

In the first eight experiments, using various discrete-trial procedures, there was a failure to obtain a significant marking effect. In Experiments 6 and 7, (and possibly Experiment 3), a counter-marking effect was found with performance in the marked group actually below that in the control group. However, a clear marking effect was found when an invisible-trials procedure was adopted in Experiments 9 to 11. This discussion is primarily aimed at trying to identify why this disparity between the results with discrete- and invisible-trial procedures might have arisen.

One possibility which is suggested by the outcomes of Experiments 9 to 11 is that marking did occur in the early experiments but that its effects were obscured by the measure of learning which was used. In Experiment 9, for instance, the performance of the marked-same and control groups as measured by percent correct (see Figure 18) did not differ substantially. However, when the number of correct responses was considered (see Figure 19), then the superiority of the marked-same group was clearly apparent. In Experiments 1 to 8 the only measure of learning used was percent correct, and that measure was based only on choice responses and not the wider data-base used in the analysis of invisible-trial experiments. In order to examine whether a marking effect emerges in the earlier studies if changes in the reponse rate are calculated rather than percent correct choices, an analysis was made of the average number of correct responses per session in Experiment 5. This experiment was chosen because it came closest to revealing a marking effect even as measured by percentage of correct choices. Unfortunately data were only available for the first half of the delay interval, but on this basis the level of correct responding was in facthigher in the control than the marked group, with 250 and 220 responses respectively. Therefore the failure to find a marking effect in the discrete-trial experiments does not seem to be an artifact of the learning measure used.

The failures in Experiments 1 to 8 seem slightly more surprising perhaps in the light of another point which emerged during the analysis of the invisible-trials procedure, where it was suggested that the benefits of marking occurred primarily on rewarded trials. In the discrete-trial studies one might expect marking to be effective on rewarded and non-rewarded trials as both food and the absence of food were clearly discriminable events at the end of a trial. If anything this should have helped produce a clearer marking effect in discrete-trial experiments.

Why then was marking found when an invisible- but not a dicrete-trial procedure was used? The difference in performance is clearly not attributable to the discrimination that subjects were required to learn. In both sets of experiments the same red-green problem was used.

One possible reason is that the red and green discriminative stimuli were present at the end of trials in the invisible-trial experiments but had been replaced with white illumination in most of the discrete-trial studies. Reinstatement theories of the kind proposed by Spear (1973, 1978) and arguments based on generalisation decrement would suggest that maintaining the same stimuli at the time of choice and reinforcement should help learning. Two points can be made about this approach. The first is that it simply predicts better learning in both marked and control group and fails to explain why markers did not facilitate the performance of marked subjects relative to controls. A second problem with this approach is that in two discrete-trial studies - Experiments 6 and 7 - the red and green cues were maintained until the end of the delay and still no significant difference between marked and control groups was obtained.

Another difference between the two sets of experiments was in the training given before the delayed-reward stage. A number of ways can be identified in which the magazine training and pretraining stages of the discrete-trial experiments might have acted to reduce the novelty of the marker and perhaps thereby devalued its effectiveness. In Experiments 1 and 2 subjects had early experience of white noise as background masking and the offset of the houselight during the food deliveries. A consequence of this experience might have been to reduce the impact of the houselight-offset marker in the first experiment and the compound of houselight-offset and white noise in the second. The same compound marker was employed in Experiment 3 and although subjects had not been exposed to white noise before the experimental phase, they had experienced the offset of the houselight many times. In Experiment 4 the group of interest is the whole-key-marker group in which choice responses were followed by the immediate darkening of the response key for 3 sec. (The half-key-marker group is not relevant as learning can be attributed to perceptual focusing in that group.) The novelty of this marker is likely to have been reduced by the procedures used in magazine training and autoshaping, both of which involved many changes in key illumination. During the pretraining stage subjects were also exposed to darkening of the key between trials.

In each of the remaining discrete-trial experiments the same magazine training, autoshaping and pretraining procedures were adopted as in Experiment 4. Therefore in so far as change in the illumination of the key was an aspect of the markers used in Experiments 5 to 8, the effectiveness of those markers might have been undermined by this experience. However, as the markers used in Experiments 5, 7 and 8 contained elements other than the darkening of the response key, this analysis offers at best only a partial explanation of the failure to obtain a marking effect. (In Experiment 6 an amber marker was used and this may have produced colour interference.)

To the extent that the interpretation offered above of the impact of early experience upon the effectiveness of the marker is correct, then the pretraining provided in the invisible-trials experiments was ideal. The marker used was a change in key illumination from red and green to white for 1 sec. Prior to the presentation of markers in delayed-reward training, the subjects had no experience of either change in the state of the response key (except for the onset and offset of the stimuli at the beginning and end of concurrent VI sessions), or white illumination of the key. The marker was therefore more likely to be effective than it had been in discrete-trial experiments.

A more general implication of the above discussion is that it would be interesting to see whether the markers used in the discrete-trial studies were effective in facilitating learning in an invisible-trial situation. Tests of this kind would not only be valuable in assessing the extent to which the markers used in Experiments 1 to 8 were responsible for the failure to obtain marking. In the case of off-the-key markers such tests would also allow the identification of events that could be more readily used in a range of further experiments on, for instance, the effect of exposure to the marker before training. A difficulty with on-the-key markers, (such as the change to 1 sec of white illumination used in Experiments 9 to 11), in, say, a test of the consequences of marker pre-exposure is that it is not possible to determine precisely the extent to which subjects have attended to them.

None of the factors considered so far has provided a comprehensive explanation for the divergence in marking between the two procedures. The familiarity of the marker might have been a factor contributing to the poor results obtained in some of the discrete-trial experiments, but as was indicated above, it cannot account for the failure to obtain marking in all of these experiments. Perhaps a more important factor is that the red and green stimuli remain on the response key throughout the session in the invisible-trial procedure but not,of course, in the discrete-trial experiments. It may be that the onset of the discriminative stimuli at the start of a discrete trial attracts the subjects' attention at a high level in both control and marked groups. If this is the case, then the attention paid by the subject to a choice response shortly afterwards might be at a ceiling such that the presentation of the marker does not promote any further attention or rehearsal. This proposal receives some support from informal observation by the experimenter that the latency to respond following the start of a trial was very short. (Unfortunately the latency of the first response was not recorded by the equipment.)

Indirect support for this view is provided by the primacy effect consistently found in studies of human short-term memory. It has been shown by Keppel and Underwood (1962) that the item presented on the first trial of an experiment is learned or remembered (or both) better than subsequent items. A primacy effect on degree of learning about the first item has also been demonstrated by Wickelgren and Norman (1966). The first response in a discrete trial might benefit similarly.

A related way in which a discrete-trial procedure might help learning is suggested by studies of short-term memory that have revealed a phenomenon called release from proactive interference (PI). (The term proactive interference is used to refer to the effects of prior learning on subsequent learning.) The effect of PI can be removed if the type of item used in a memory experiment is changed after a series of items of one kind. Memory for the first item of a new category is much improved (Loess, 1968; Wickens, Born and Allen, 1963). A substantial release from PI has also been found simply by waiting for a short time between memory trials (Peterson and Gentile, 1965; Loess and Waugh, 1967; Kincaid and Wickens, 1970). Thus in a discrete-trial situation the onset of discriminative stimuli may help all subjects (even if a choice response does not necessarily follow shortly) by breaking up responses into separate groups, the first one of which is learned about well. The marker, therefore, may not be able to provide any further enhancement of learning.

In contrast with the discrete-trial procedure, in invisible-trial experiments the choice response is not preceded by any change in key illumination and, as far as the subject is concerned, will simply be one in a very long series of pecks. One might therefore expect the presentation of a marker to have much more effect upon the subjects' attention.

One way to test this hypothesis would be to vary the duration of the period between the onset of the discriminative stimuli and the marked response. If this analysis of the difference between the level of marking in the two procedures is correct, then one would anticipate that the longer the period of exposure to the stimuli before a response is marked, the greater the facilitation of learning by the marker.

An alternative, or complementary, factor might be the number of responses that occur before a marked response. Discrete-trial marking may be less effective because a marker following the first response of a series may be less surprising than one which follows a number of non-marked responses. The latter situation will prevail with the invisible-trials procedure.

In conclusion, the pretaining procedures used in the discrete-trial experiments might have devalued the markers subsequently used whereas no such problem appeared possible in the invisible-trial experiments. However, the major factor that might have removed any benefit of marking in the experiments 1 to 8 is the attention-attracting properties of the onset of the discriminative stimuli at the start of each trial. This might have significantly assisted learning in the controls and thus made it substantially harder to uncover any facilitating effect of marking.

5.5 VISUAL AND SPATIAL MARKING

A pattern which has emerged from studies of marking, including those reported here, suggests that it may be more difficult to obtain a marking effect with visual discriminations than spatial ones. In order to provide a basis for an examination of why this divergence between visual and spatial marking might occur, the relevant evidence is first reviewed.

The experiments performed to investigate marking (or which have involved designs susceptible to a marking analysis) fall into three categories. The first is composed of studies with rats in mazes, the second with pigeons in operant chambers using discrete-trial procedures, and the third also with pigeons in operant chambers but using an invisible-trials design.

To begin with the first category. Lett (1973, 1975) has reported that rats can learn a spatial discrimination in a T-maze with delays as long as 1 hour. Lett marked her subjects by handling following choice responses. Using a similar procedure to Lett (1973), Lieberman et al. (1979, Experiments 1 and 2) confirmed that marked rats could learn a spatial discrimination with delays of 1 min. Five further experiments by Lieberman et al. (1979) and Thomas et al. (1983) also showed learning over delays of 30 sec, 1 min and 2 min in several different mazes (see Figure 1 for the basic layout). In total, these experiments suggest that spatial marking is a highly reliable phenomenon with rats in mazes.

The picture with visual marking is much less consistent. Lett (1974) reported that rats could learn a visual discrimination in a T-maze with a 1-min delay of reward. In Lett (1978), she further reported learning with delays ranging from 0 to 240 sec. However, doubt was cast on the reliability of the effect by Roberts (1976) who reported a failure to replicate Lett's (1974) finding of learning with a 1-min delay of reinforcement. Subsequently, Lett (1977) reanalysed Roberts' (1976) data and claimed that in fact his results did show learning in the group that precisely repeated her original design. In his reply to Lett, Roberts (1977) argued - convincingly in my opinion - that she had used a questionable statistical procedure to reanalyse his results and that anyway, if she was correct, then it simply showed that the learning she had reported arose from the use of odour cues. The essential conclusion from this debate for the present discussion is that handling after choice responses does not reliably facilitate visual discrimination learning by rats. In the three undisputed groups contained in Roberts' (1976) study, no learning was found despite marking of responses.

Further doubt was cast on the reliability of visual marking by Young (1979) and Coward (1980). In one experiment using the same basic maze as in Thomas et al. (1983) and a 1.5-sec burst of white noise as a marker, Young found that rats could learn a black-white discrimination despite a 30-sec delay of reinforcement. However, he was unable to find learning in either marked or control groups with a discrimination between horizontal- and vertical-black stripes on a white background with the same delay. Coward (1980) then attempted to replicate Young's first study very closely but found no learning in either marked or control groups.[1]

---

[1] The difference in the marked-group results between Young's (1979) first experiment and Coward's (1980) study might in part be accounted for by the fact that Young used two markers - one after the choice response and the second at the end of the delay - whereas Coward only followed choice responses with a marker. A potential difficulty with this analysis is that the control group showed an improvement in Young's study but not in Coward's. However, inspection of the learning curves of the two control groups suggests that the initial colour preference was much stronger in Young's subjects. Therefore their change in percent correct, (from about 10% to 40% over the course of training), might simply reflect the loss of this preference rather than learning about which colour is correct.

doubt was cast on the reliability of the effect by Roberts (1976) who reported a failure to replicate Lett's (1974) finding of learning with a 1-min delay of reinforcement. Subsequently, Lett (1977) reanalysed Roberts' (1976) data and claimed that in fact his results did show learning in the group that precisely repeated her original design. In his reply to Lett, Roberts (1977) argued - convincingly in my opinion - that she had used a questionable statistical procedure to reanalyse his results and that anyway, if she was correct, then it simply showed that the learning she had reported arose from the use of odour cues. The essential conclusion from this debate for the present discussion is that handling after choice responses does not reliably facilitate visual discrimination learning by rats. In the three undisputed groups contained in Roberts' (1976) study, no learning was found despite marking of responses.

Further doubt was cast on the reliability of visual marking by Young (1979) and Coward (1980). In one experiment using the same basic maze as in Thomas et al. (1983) and a 1.5-sec burst of white noise as a marker, Young found that rats could learn a black-white discrimination despite a 30-sec delay of reinforcement. However, he was unable to find learning in either marked or control groups with a discrimination between horizontal- and vertical-black stripes on a white background with the same delay. Coward (1980) then attempted to replicate Young's first study very closely but found no learning in either marked or control groups.[1]

---

[1] The difference in the marked-group results between Young's (1979) first experiment and Coward's (1980) study might in part be accounted for by the fact that Young used two markers - one after the choice response and the second at the end of the delay - whereas Coward only followed choice responses with a marker. A potential difficulty with this analysis is that the control group showed an improvement in Young's study but not in Coward's. However, inspection of the learning curves of the two control groups suggests that the initial colour preference was much stronger in Young's subjects. Therefore their change in percent correct, (from about 10% to 40% over the course of training), might simply reflect the loss of this preference rather than learning about which colour is correct.

The second category of experiments is those peformed with pigeons using a discrete-trial procedure. In Experiments 1 to 8 of this thesis, no significant marking effect was found even though in all but one experiment the task was clearly capable of solution as learning occurred. On the other hand with a spatial discrimination and essentially the same design, Fuller (1981) and Lieberman (personal communication, 1981) found better learning in their marked subjects than their controls.

The final category of experiments involves pigeons trained with an invisible-trials procedure. As the general discussion at the end of Chapter 4 indicated, there appeared to be a stronger marking effect in Lieberman's (personal communications, September 1983, January 1984) spatial experiments than in Experiments 9 and 11.

A part of the reason for the disparity between spatial and visual situations might lie with the relative difficulty of the two types of problems. Mackintosh (1974, p.604) reported that spatial tasks are typically easier than visual ones, a view confirmed by Bitterman (1965) and Bullock and Bitterman (1962). If Roberts (1976, 1977) is correct, then there really is little evidence of learning, let alone marking, in visual discrimination problems with rats over long delays. The first Young (1979) experiment might represent a rare demonstration of learning in an adequately controlled design. However, the relative difficulty of visual and spatial tasks is only adequate as an explanation for the divergence in marking effects in so far as the visual problems used are too difficult for any learning. Discriminations of that kind may be such that even following choice responses with salient events does not promote solution of the problem. However, in so far as learning occurs, (as it clearly does in the pigeon experiments), then some other factor is required to explain the apparently stronger spatial marking effect.

One possibility is that the markers used in the spatial experiments were more salient or novel. Whilst this might have been true for the pigeon, invisible-trial experiments, (where a 1-sec flash of red-green key illumination was used in the spatial studies and a flash of white illumination in the visual task), it is clearly not applicable in the rat experiments where handling was used in both spatial and visual situations.

An alternative account which was proposed in the conclusion to Chapter 3 is that subjects might generally have coded problems initially in spatial terms. By the time they came to attend to the relevant dimension in visual discriminations, the effectiveness of the marker might have been reduced as a result of repeated presentations. Analysis of the pretraining behaviour of the pigeons, however, did not support this interpretation as the subjects displayed visual as well as spatial preferences from the outset. A review of the literature on visual discrimination learning in rats by Mackintosh (1974, pp.571-573) suggests that rats usually initially respond on the basis of position. But he concluded that some learning does occur to visual stimuli even while animals are responding to position. This implies that marking should still proceed during the early stages of visual discrimination learning in rats, but that it may be less effective than in the initial stages in spatial tasks. Thus, if repeated exposure to the marker does reduce its effectiveness, then one might expect to find a somewhat weaker marking effect in visual than spatial problems when rats are used as subjects.

A more powerful general explanation of the divergence between the two types of discrimination is suggested by a consideration of the cues available to the subjects. In all of the visual discrimination experiments irrelevant spatial cues were necessarily present. Marking might therefore have facilitated learning about irrelevant spatial cues instead of, or as well as, about relevant visual cues. In contrast, in all of the spatial discrimination tasks discussed above there were either no salient visual cues, as in the split-key apparatus used by Fuller (1981) and Lieberman (personal communications, 1981, September 1983, January 1984) in which the halves of the key were lit white, or the visual cues were redundant relevant stimuli (Lett, 1973, 1975; Lieberman et al., 1979; Thomas et al., 1983). In the latter experiments marking of visual cues would not impair learning as they were relevant to the discrimination. In Lieberman et al. (1979), for example, the left side-arm of the maze was always black and the right always white.

This discussion suggests that more experiments need to be conducted before it is concluded that the spatial marking effect is inherently stronger than the visual marking effect. In particular, it would be informative to introduce irrelevant visual cues into spatial tasks in order to allow more controlled comparisons of visual and spatial marking. For example, irrelevant red and green stimuli could be added to the designs used by Lieberman (personal communications, September 1983, January 1984) and the results compared with those obtained in Experiments 9 and 11. Incidentally, a change of this kind would allow the same marker to be used in both visual and spatial versions thereby allowing a highly controlled comparison to be made.

5.6 MARKING AND WAGNER'S REHEARSAL MODEL

An interesting issue raised by marking is its relationship to current models of information-processing in conditioning. One such theory that has been highly influential is Wagner's (1976, 1978, 1979, 1981) model of animal short-term memory (STM) and his account of the role of STM in the formation of associations. His model includes the following assumptions.

1. Short-term memory has a limited capacity.
2. Permanent associative connections are developed between representative elements only to the degree that the elements are jointly active in short-term memory.
3. Events are maintained in memory through a rehearsal or memory-maintenance process.
4. Surprising or unexpected events are subjected to different degrees or different "levels" of processing than are expected events. Surprising events are more likely to be "rehearsed" or maintained in a primary, A1, state, than are expected events.
5. In engaging the rehearsal mechanism, surprising events take up limited capacity short-term memory and displace other representations.

Wagner also defines a surprising or unexpected event as one that is not predicted by the total aggregate of ones that precede it. Expected events can be prerepresented in short-term memory as a result of the retrieval action of other antedating stimuli that they are associated with, or by the recent presentation of the same event.

How do the various elements of the marking situation fit into the framework provided by Wagner? His model has been framed in terms of classical conditioning and the way in which CSs and USs become associated together. However, there seems to be no good reason why responses should not be represented in STM in the same way as stimuli. If this extension is accepted, then Wagner's model suggests that a response activates a representative mode, or set of memory elements, just as stimuli do (Wagner, 1981, p.9).

It is clear that Wagner's definition of a surprising event - one that is not predicted by the total aggregate of ones which precede it - encompasses markers. The markers used in Lieberman et al's (1979) experiments - handling, white noise, and a light flash - were not events that the choice response or the side arms of the maze predicted. In particular, those markers which were least familiar and certainly least expected in the context of the maze, (the white noise and light flash), were the most effective. Similarly, the change in key illumination from red and green to white used as a marker in Experiments 9 and 11 was highly novel. Before the training phase subjects had very limited experience of changes in the state of the key and had never seen it lit with white light.

Given that markers are surprising events in Wagner's terms, how might they be expected to operate? Thomas et al. (1983) took the view that "the model predicts that a marker should hamper learning because it will interfere with rehearsal of the preceding response". This interference effect will occur because rehearsal capacity is limited, so that any processing given to a surprising event such as a marker will reduce the processing capacity available for the preceding response, and thereby lower the probability of the response being remembered. In the language of Wagner's (1981) most sophisticated version of the model, the marker will displace some elements of the response mode from the primary (A1) state of activity in STM. In control subjects there is no marker to produce such a displacement and the response node will remain in the primary state for longer.

Though the memorial implications of Wagner's model for marking seem quite clear, the picture concerning the formation of associations is less obvious. Two predictions do seem to follow with some certainty from the 1981, SOP version of the model.

1. The marker enters into an association with choice responses because there is conjoint processing of the A1 state of the choice response node with the A1 state of the marker node.
2. The marker will enter into associations with both correct and incorrect choices as it follows all choice responses.

A number of other deductions from the model may be possible and some of these are listed below.

1. The excitatory association formed between the correct response and reinforcement will be less in the marked condition than the control one because the marker will reduce the level of A1 activity produced by the correct choice, as explained above.
2. On correct choice trials, the marker will, by virtue of its association with the incorrect choice, promote the incorrect choice node into a state of secondary, A2, activity. To the extent that this activity continues until the reinforcer is delivered, an inhibitory association will be formed between the incorrect choice and the reinforcer. In the control condition, there is no marker to evoke a representation of the incorrect choice on correct trials, and therefore inhibitory learning will be less.

3. The marker may enter into an excitatory association with the reinforcer to the extent that elements of its node are still in an A1 state by the time the reinforcer is presented (though this also depends upon the level of marker A2 activity at the end of the delay which determines inhibitory conditioning).
4. If the marker does form an excitatory association with the reinforcer, then it will "prime" some of the elements of the reinforcer mode into an A2 state (Wagner, 1981, pp.11 and 39) thereby possibly reducing the level of A1 activiation when the reinforcer itself is presented.

What is obvious from this discussion is that it is very difficult to determine what the net effect of the marker will be in associative terms. Much depends upon the precise parameters that are, one, adopted to describe the probability with which elements enter into particular states of activation and, two, used to determine the rate of decay of those elements.

It is impossible at this time to derive any predictions from Wagner's model about the relationship between the marked and control conditions in terms of the formation of associations. It does, however, seem apparent that a marker should be damaging to memory for the response according to the model. In order to reconcile this with the facilitation of learning found in marking, Thomas et al. (1983, p.410) propose a possible change to the rehearsal model. They suggest that

> "the reduced amount of available processing is not necessarily redistributed equally over the set of events concerned. In particular, if an event with some special relationship to the salient stimulus (e.g., contiguity, similarity) receives a larger share of the reduced amount of processing, then this event could actually receive more processing rather than less following a salient stimulus".

An amendment of this kind might involve the kind of controlled processing routines, mentioned by Wagner (1981, p.43), that can produce a transition of selected memory nodes from an A2 to an A1 state by an act of attention.

5.7 MARKING AND RESPONSE OVERSHADOWING

In marking, following a response with a brief stimulus facilitates the learning of an association between that response and reinforcement. However, in some other situations the addition of a brief stimulus interferes with learning. One such effect has been reported by Pearce and Hall (1978) who investigated the effect of giving rats an occasional brief response-contingent stimulus while lever pressing for food on a VI schedule. They found that if the stimulus occurred only on rewarded responses, then the response rate was lowered relative to a condition in which the stimulus was presented just as often but uncorrelated with reinforcement. Lever pressing in the uncorrelated condition was at the same level as in control subjects which simply received reinforcement on the same schedule but without the response-contingent stimulus. A similar finding was reported by St. Claire-Smith (1979a) who also demonstrated an analogous effect with rats punished for level pressing on a fixed-interval (FI) 5-min schedule (St. Claire-Smith, 1979b). The generality of the effect is further illustrated by Hall's (1982) finding that a response-contingent stimulus correlated with reinforcement was effective in lowering the rate of food-reinforced pecking in pigeons. And even more recently Tarpy, Roberts, Lea and Midgley (1984) reported evidence suggesting that the same effect is found with rats rewarded on an FI schedule.

Pearce and Hall (1978) interpreted their results in the context of the proposal (Bolles, 1972; Mackintosh, 1974, pp.206-207) that instrumental learning involves the formation of response-stimulus associations. They proposed that the effect was analogous to a finding frequently obtained in classical conditioning. Wagner (1969), for example, found that if two stimuli are reinforced in compound but one is made a less reliable or valid predictor of reinforcement, then less conditioning occurs to that stimulus than if it had received the same reinforcement correlation on its own. The more valid one is said to "overshadow" the less valid one. Pearce and Hall (1978) argued that the lower rate of responding found when the brief stimulus was correlated with reinforcement could be attributed to an overshadowing effect. The stimulus was always followed by food and therefore became a better predictor of reinforcement than responses, many of which were not reinforced. The stimulus-reinforcer association was said to overshadow the response-reinforcer association.

An analysis in terms of the relative predictiveness of stimuli is capable of explaining why the interference effects of brief stimuli found by Pearce and Hall (1978) and others did not occur in the studies carried out by Lieberman et al. (1979) and Thomas et al. (1983). In the former experiments the stimulus is a better predictor of reinforcement than the response it followed. On the other hand, in the marking experiments by Lieberman and his colleagues, a correct response is a more valid predictor of reinforcement. Correct choices are always reinforced whereas markers follow both correct and incorrect choices.

However, this account of why response-contingent stimuli interfere with learning in some studies but not in others is not supported by the results of invisible-trial marking experiments (Experiments 9 and 11; Lieberman, personal communications, September 1983, January 1984). In these experiments the marker was a much more reliable predictor of reinforcement than a correct response, but nonetheless learning was facilitated in marked-same groups relative to controls.

An alternative to the relative validity account discussed above is suggested by considering the temporal relationship between response, stimulus and reinforcer. In all of the marking experiments there has been a delay between the response and reinforcement. The duration of the delay ranged from 6 sec (Experiments 5 to 8; Fuller, 1981; Lieberman, personal communication, 1981) to 2 min (Lieberman et al., 1979, Experiment 4). In contrast, in the work by, for example, Pearce and Hall (1978) reinforcement was either immediate or delayed by only 0.5 sec.

A way in which this difference in temporal relationships might affect the processing of events is suggested by Wagner's (1978, 1981) analysis of memory processing in classical conditioning. When a US is first presented, its initial effect is assumed to be to reduce processing of the CS, but as their representations remain simultaneously active in memory an association is strengthened between them. If a surprising event is presented soon after the trial, however, this post-trial episode will reduce the simultaneous processing of the CS and US and thus reduce learning (Wagner, 1981, pp.38-39).

Applying this analysis to marking in the modified form proposed by Thomas et al. (1983), when the marker is first presented, one might expect an initial reduction in processing of the response but with continued time for processing the marker may initiate a memory search which ultimately results in enhanced response processing. If reinforcement follows the marker too soon, however, there will not be time for the response rehearsal phase to have been initiated, so that the main effect of the marker will simply be to reduce processing of the response, and thereby reduce association of the response with reinforcement.

In summary then, Pearce and Hall's (1978) overshadowing results and marking can be reconciled if it is assumed that an immediate effect of the marker is to reduce processing of the response (as suggested by Wagner's model), but that with further time processing of the response is actually increased.

However, the generality of the response overshadowing effect has recently been challenged by the finding that it does not occur with VR schedules. Roberts, Tarpy and Lea (1984) and Tarpy, Lea and Midgley (1983) found no difference between a condition where the response-contingent stimulus was correlated with food on a VR schedule and one in which the stimulus randomly followed responses. Neither relative validity nor processing time can account for the difference between effects obtained using VI and VR schedules. In both schedules the relative validity of the response and the stimulus as predictors of reinforcement is the same. Similarly, the temporal relationship between response, stimulus and reinforcer is identical in each schedule.

In order to explain the discrepancy between VI and VR, Roberts et al. (1984, p.254) have suggested that the correlated stimulus "highlights the reinforcement contingencies by making the temporal interval offset on an interval schedule of reward more salient".[1]

---

[1] Roberts et al. (1984) also interpret their results as indicating that subjects in the correlated-stimulus condition had developed "a better knowledge of the two causality sources of reward (time elapse and responding)". They claim that this knowledge is reflected in the lower response rate by correlated subjects, which economise on their expenditure of energy without sacrificing rewards. Although this analysis might make sense in the context of food reinforcement, it seems difficult to apply it to the punishment results reported by St. Claire-Smith (1979b). He found that a stimulus preferentially correlated with shock reduced the extent to which punishment suppressed the rate of responding. If one effect of the stimulus was to improve the subjects' knowledge of the contingencies of reinforcement as Roberts et al. (1984) claim, it seems unlikely that subjects punished for lever pressing would continue to do so at a higher rate than those for which shock is not preceded by a brief stimulus.

In some ways an interpretation of this kind is attractive from a marking perspective as it can be viewed as an extension of marking rather than a challenge to it. The process responsible for making the interval more salient could be marking. It should be noted, however, that the temporal-interval salience interpretation cannot explain the invisible-trials marking results. Although application of the analysis to the invisible-trials procedure is complicated by the fact that the marker is not completely correlated with food, the marker should nonetheless still enhance the salience of temporal interval offset. One might therefore expect a lowering of response rate relative to a control condition without markers. Moreover, whatever predictions are derived from Roberts et al.'s (1984) analysis, they should presumably have the same implications for the marked-same and marked-different groups. But there is clear evidence from Experiments 9 and 11 that, one, the marker does not reduce response rate and, two, the behaviour of the marked-same and marked-different groups differs.[1] Therefore marking of responses clearly occurs and if marking of time does take place, then it is a supplementary process.

The picture that emerges from this survey of studies of the effect of following responses with brief stimuli, is complicated. In some situations such stimuli have been found to facilitate learning (for example, Experiments 9 and 11), in others to make no difference to performance (Tarpy et al., 1983, Experiments 1, 3 and 4; Roberts et al., 1984, Experiment 3), and in a third group, to interfere with

---

[1] It might be interesting to examine the effects of adding additional response-contingent stimuli to the correlated VI condition of the overshadowing experiments in such a way as to make the condition more similar to the invisible-trials marking procedure and thereby allow a more direct comparison.

learning (Hall, 1982; Pearce and Hall, 1978; Roberts et al., 1984, Experiments 1 and 2; St. Claire-Smith, 1979a, 1979b; Tarpy et al., 1983, Experiment 2, 1984). Neither the response overshadowing analysis offered by Pearce and Hall (1978) to account for their results, nor the temporal interval salience alternative proposed by Roberts et al. (1984) can explain the facilitation of learning found in Experiments 9 and 11. The marking hypothesis appears to offer the only satisfactory explanation for this enhancement of learning.

On the other hand, if marking is to account for the results presented by Pearce and Hall (1978), etc. and by Tarpy and his colleagues, then the hypothesis needs to be supplemented. Perhaps this could be done by adding an assumption that marking requires time without the presentation of other events such as a reinforcer. An additional possibility is that marking of time intervals can occur as well as responses.

In order to obtain a more complete picture of the effects of brief stimuli upon learning and, one hopes, a better understanding of the processes involved, it would be valuable to investigate the processing time and relative validity factors using designs like those in the marking studies. The relative validity of the marker and the response could be varied, as could the timing of the response-marker-reinforcer sequence.

5.8 MARKING AND DELAYED MATCHING-TO-SAMPLE

Another area in which the addition of stimuli has been found to interfere with learning is delayed matching-to-sample (DMTS) in both pigeons (Cook, 1980; Grant and Roberts, 1976; Maki, 1979; Roberts and Grant, 1978; Tranberg and Rilling, 1980) and monkeys (D'Amato and O'Neill, 1971; Worsham and D'Amato, 1973). In delayed matching with pigeons a three-key chamber is normally used, and the sample stimulus is presented on the centre key. After the sample has been available for a few seconds, the key is darkened and a delay interval follows. The comparison stimuli are then presented on the side keys and a response to the stimulus that matches the sample is reinforced. If an additional stimulus is interpolated during the delay interval, then the accuracy of matching performance is reduced.

A number of researchers have suggested that the interference effects they have found are at least consistent with Wagner's (1976, 1978, 1981) rehearsal model (Cook, 1980; Grant and Roberts, 1976; Roberts, 1980; Roberts and Grant, 1978). The interpolated stimulus is seen as disrupting the rehearsal of the sample and displacing it from a limited-capacity short-term memory. Some doubts can be raised about the extent to which Wagner's model completely fits the data - for example, Grant and Roberts (1976) found that novel interpolated stimuli were no more disruptive than familiar ones (see also, Tranberg and Rilling, 1980) - but nonetheless, the results are broadly in line with the model.

The interference effect of the additional stimuli in DMTS is contrary to what one might have expected in the light of the marking hypothesis. A marking analysis might predict that the interpolated stimulus at least in so far as it follows the sample, immediately would initiate a backward search through memory. This search would identify the sample stimulus as the predictive cue and therefore mark it. Thus, when the comparison stimuli are presented, the sample stimulus should be more memorable than in a control condition where the centre key was simply darkened during the delay interval. (Although the darkening of the key following a sample stimulus might also be expected to generate a backward search, one would anticipate a stronger effect with the more salient stimuli used as interpolated events). One way to approach this disparity between facilitatory predictions which can be derived from a marking analysis and the interference effect obtained, is to examine the DMTS procedure carefully in order to identify potentially important differences to marking experiments. To avoid the complications of cross-species comparisons this discussion will focus on DMTS and marking experiments that have used pigeons as subjects.

One general difference between the two types of studies is in the duration of the additional stimulus. In Experiments 9 to 11, for example, the marker was a 1-sec change in key illumination, whereas in most DMTS work the interpolated stimulus filled the entire delay interval. A brief "flash" of a marker may be more effective. However, that this is not a crucial difference is shown by Roberts and Grant (1978, Experiment 2) who varied the length of the interpolated stimulus from 1 to 8 sec. Although the level of a disruption was significantly lower with the 1 sec stimulus, the effect was nonetheless one of interference.

A second possibility is that the type of stimulus used determines the nature of its effect. The invisible-trials marking experiments reported in this thesis involved a change in key illumination to white for 1 sec. In the majority of DMTS studies the onset or offset of the houselight has been used as the interpolated event. However, two pieces of evidence suggest that this difference is not responsible for the divergence in effects. The first is that when houselight offset was used as a marker in Experiments 5 and 6, the marked groups' performance was slightly (though not significantly) better than the control groups'. No interference effect was found. The second result of importance for the present comparison is Grant and Roberts (1976) finding that interference is obtained when a white-on-black pattern is used as the interpolated stimulus. Although this pattern is clearly not the same as the white marker adopted in Experiments 9 to 11, these results certainly strongly suggest that the type of stimulus used in DMTS and marking experiments is not responsible for the different effects found.

A third alternative is the two sets of studies involve a crucial dissimilarity of procedure in terms of the response keys that are used. The marking experiments reported here all involved a single key which was split in half. The majority of DMTS studies used the centre key for the presentation of the sample and the side keys for the comparison stimuli. This difference might be significant in some way. This interpretation, however, is not supported by the results of an experiment by Tranberg and Rilling (1980) employing a successive DMTS task. Both sample and comparison stimuli were presented on the same key. The sample was given first, a delay followed, and then a single comparison stimulus was presented. Reinforcement was made available after the comparison if it matched the sample, but not if it was different from the sample. Responding during matching and non-matching comparison stimuli relative to baseline responding was used to assess performance. Tranberg and Rilling (1980) found interference effects using this successive DMTS procedure similar to those obtained with the more popular, two-choice DMTS task.

A fourth, and possibly more important difference is that in the marking studies choice responses were followed by an additional stimulus, whereas in all but one of the pigeon DMTS experiments mentioned above the interpolated event followed the sample stimulus. Thus, although chance pairings of a response to sample and a change to the interpolated stimulus might have occurred, this contingency was not programmed. It may be that for some reason marking of responses is a more powerful process than marking of stimuli. Indirect support for this view might be derived from Grant's (1981) proposal that in DMTS (and other areas) the subject encodes in memory an instruction to, say, "peck green" rather than the physical characteristics of the sample stimulus. However, a difficulty for this analysis in terms of the relative ease of marking stimuli and responses is raised by Tranberg and Rilling's (1980) study. In their experiment the length of presentation of sample stimulus was determined on a FI 5-sec schedule. Thus, the presentation of the interpolated event always followed immediately after a peck to the sample. Nonetheless interference was still found.

Although an analysis based on the relative effectiveness of marking stimuli and responses is not supported by the Tranberg and Rilling (1980) experiment, it does suggest a way in which many of the DMTS studies might have been made more difficult by an interpolated stimulus. In most DMTS procedures the first event of a trial was the illumination of the centre key with white light. A response on a fixed-ratio (FR) 1 schedule produced the sample, the onset of which might have marked the peck to the white stimulus. Later when the sample was replaced by the interpolated stimulus a second scan through memory might have been initiated which might have identified the marked response to the white key as the predictive cue rather than the sample stimulus. In so far as the interpolated event produced a more effective memory scan than the simple darkening of the key (as in the control condition), then subjects experiencing an interpolated event might have been more likely to identify erroneously a peck to white than control subjects.

A final factor which might be of significance is suggested by the earlier discussion of the difference between discrete-trial and invisible-trial marking experiments. It was suggested that marking may be difficult with a discrete-trial design because of the attention-attracting properties of discriminative stimulus onset. The DMTS studies clearly involve a situation closer to that found in discrete-trial rather than in invisible-trial experiments. It may be that interference effects are more likely to occur when the attention of the subject is already at a high level.

What is apparent from this discussion is that it is not easy to identify those factors responsible for the finding that additional stimuli can sometimes facilitate learning and in other situations interfere with it. It may be useful to modify the DMTS procedure in ways which this analysis has suggested would allow a more direct comparison with marking studies. For instance, the event successfully used as a marker in Experiments 9 to 11 could be adopted as the interpolated event in a DMTS procedure. Another possibility is that the termination of the sample stimulus in DMTS could be made response contingent and also that the sample is no longer preceded by white illumination. However the general picture is one which suggests that a distractor effect is dominant in DMTS situations. It is worth noting that other investigators have concluded that the processes involved in discrimination learning and delayed matching may be different (Maki, 1979; Roberts, 1980, p.229).

This examination of the relationship between DMTS and marking, and the earlier discussion of response overshadowing make clear that the field of animal STM is a complicated one. Other areas of research such as potentiation (Holland, 1980; Palmerino, Rusiniak and Garcia, 1980; Pearce, Nicholas and Dickinson, 1981) and the effects of interpolating a stimulus between a CS and US (Rescorla, 1980) will also require integration into a model of animal STM. What is apparent from the experiments reported here is that an additional stimulus can facilitate learning in some circumstances and that an adequate model will have to incorporate this phenomenon as well as the interference effects uncovered elsewhere.

REFERENCES

BADDELEY, A.D. (1976). The psychology of memory. New York: Harper & Row.

BAUM, W.M. (1973). The correlation-based law of effect. Journal of the Experimental Analysis of Behaviour, 20, 137-153.

BERLYNE, D.E. (1967). Arousal and reinforcement. In D. Levine (Ed.) Nebraska Symposium on Motivation 1967, pp. 1-110. Lincoln: University of Nebraska Press.

BERLYNE, D.E., BORSA, D.M., HAMACHER, J.H., and KOENIG, I.D. (1966). Paired-associate learning and the timing of arousal, Journal of Experimental Psychology, 72, 1-6.

BITTERMAN, M.E. (1965). Phyletic differences in learning. American Psychologist, 20, 396-410.

BULLOCK, D.H. and BITTERMAN, M.E. (1962). Habit reversal in the pigeon. Journal of Comparative and Physiological Psychology, 55, 958-962.

CARTER, D.E. and WERNER, T.J. (1978). Complex learning and information processing by pigeons: A critical analysis. Journal of the Experimental Analysis of Behaviour, 29, 565-601.

CATANIA, A.C. (1962). Independence of concurrent responding maintained by interval schedules of reinforcement. Journal of the Experimental Analysis of Behaviour, 5, 175-184.

CATANIA, A.C. (1966). Concurrent operants. In W.K. Honig (Ed.), Operant behaviour: Areas of research and application. Englewood Cliffs, NJ : Prentice-Hall.

CATANIA, A.C. and CUTTS, D. (1963). Experimental control of superstitious responding in humans. Journal of the Experimental Analysis of Behaviour, 6, 203-208.

COOK, R.G. (1980). Retroactive interference in pigeon short-term memory by a reduction in ambient illumination. Journal of Experimental Psychology: Animal Behaviour Processes, 6, 326-338.

COWARD, C. (1980). Visual discrimination learning when reward is delayed: A further investigation of the marking hypothesis. Unpublished B.Sc. dissertation, Stiring University.

CRONIN, P.B. (1980). Reinstatement of postresponse stimuli prior to reward in delayed-reward discrimination learning by pigeons. Animal Learning and Behaviour, 8, 352-358.

D'AMATO, M.R. and O'NEILL, W. (1971). Effect of delay-interval illumination on matching behaviour in the capuchin monkey. Journal of the Experimental Analysis of Behaviour, 15, 327-333.

D'AMATO, M.R., FAZZARO, J. and ETKIN, M. (1968). Anticipatory responding and avoidance discrimination as factors in avoidance conditioning. Journal of Experimental Psychology, 77, 41-47.

DAVIDSON, F. (1983). The role of marking when reward is delayed: The "marker as a reward" hypothesis. Unpublished B.Sc. dissertation, Stirling University.

DE VILLIERS, P.A. (1977). Choice in concurrent schedules and a quantitative formulation of the law of effect. In W.K. Honig and J.E.R. Staddon (Eds.), Handbook of operant behaviour (pp. 233-87). Englewood Cliffs, NJ: Prentice-Hall.

DEWS, P.B. (1960). Free-operant behaviour under conditions of delayed reinforcement. 1. CRF-type schedules. Journal of the Experimental Analysis of Behaviour, 3, 221-234.

DOTY, B.A. and DOTY, L.A. (1966). Facilitative effects of amphetamine on avoidance conditioning in relation to age and problem difficulty. Psychopharmacologia, 9, 234-341.

EASTERBROOK, J.A. (1959). The effect of emotion on cue utilization and the organisation of behaviour. Psychological Review, 66, 183-201.

FETTERMAN, J.G. and STUBBS, D.A. (1982). Matching, maximizing, and the behavioural unit: Concurrent reinforcement of response sequences. Journal of the Experimental Analysis of Behaviour, 37, 97-114.

FOWLER, H. and WISCHNER, G.J. (1969). The varied functions of punishment in discrimination learning. In B.A. Campbell and R.M. Church (Eds.), Punishment and aversive punishment (pp.375-420). New York: Appleton-Century-Crofts.

FULLER, A. (1981). Spatial discrimination learning when reward is delayed: further research towards understanding marking hypothesis. Unpublished B.Sc. dissertation, Stirling University.

GRANT, D.S. (1976). Effect of sample presentation time on long-delay matching in the pigeon. Learning and Motivation, 7, 580-590.

GRANT, D.S. (1981). Short-term memory in the pigeon. In N.E. Spear and R.R. Miller (Eds.), Information processing in animals: Memory mechanisms. Hillsdale, NJ: Erlbaum.

GRANT, D.S. and ROBERTS, W.A. (1976). Sources of retroactive inhibition in pigeon short-term memory. Journal of Experimental Psychology: Animal Behaviour Processes, 2, 1-16.

GRAYSON, R.J. and WASSERMAN, E.A. (1979). Conditioning of two-response patterns of key pecking in pigeons. Journal of the Experimental Analysis of Behaviour, 31, 23-29.

GREEN, R.T. (1956). Surprise as a factor in the von Restorff effect. Journal of Experimental Psychology, 52, 340-344.

GREEN, R.T. (1958). Surprise, isolation and structural change as factors affecting recall of a temporal series. British Journal of Psychology, 49, 21-30.

HALL, G. (1982). Effects of a brief stimulus accompanying reinforcement on instrumental responding in pigeons. Learning and Motivation, 13, 26-43.

HAMILTON, P., HOCKEY, G. and QUINN, J. (1972). Information selection, arousal and memory. British Journal of Psychology, 63, 181-189.

HERRNSTEIN, R.J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. Journal of the Experimental Analysis of Behaviour, 4, 267-272.

HERRNSTEIN, R.J. (1970). On the law of effect. Journal of the Experimental Analysis of Behaviour, 13, 243-266.

HOLLAND, P.C. (1983). Representation-mediated overshadowing and potentiation of conditioning aversions. Journal of Experimental Psychology: Animal Behaviour Processes, 9, 1-13.

JENKINS, W.O. and POSTMAN, L. (1948). Isolation and spread of effect in serial learning. American Journal of Psychology, 61, 214-221.

JONES, F.N. and JONES, M.H. (1942). Vividness as a factor in learning lists of nonsense items. American Journal of Psychology, 55, 96-101.

KAMIN, L.J. (1968). 'Attention-like' processes in classical conditioning. In M.R. Jones (Ed.), Miami symposium on the prediction of behaviour: Aversive stimulation (pp.9-33). Coral Gables, Fla: University of Miami Press.

KAMIN, L.J. (1969). Predictability, surprise, attention and conditioning. In B.A. Campbell and R.M. Church (Eds.), Punishment and aversive behaviour. New York: Appleton-Century-Crofts.

KELLEHER, R.T. (1961). Schedules of conditioned reinforcement during experimental extinction. Journal of the Experimental Analysis of Behaviour, 4, 1-5.

KEPPEL, G. and UNDERWOOD, B.J. (1962). Proactive inhibition in short-term retention of single items. Journal of Verbal Learning and Verbal Behaviour, 1, 153-161.

KINCAID, J.P. and WICKENS, D.D. (1970). Temporal gradient of release from proactive inhibition. Journal of Experimental Psychology, 86, 313-316.

KISH, G.B. (1966). Studies of sensory reinforcement. In W.K. Honig (Ed.), Operant behaviour: Areas of research and application, (pp. 109-159). New York: Appleton-Century-Crofts.

KLEINSMITH, L.J. and KAPLAN, S. (1963). Paired-associate learning as a function of arousal and interpolated interval. Journal of Experimental Psychology, 65, 190-193.

KLEINSMITH, L.J. and KAPLAN, S. (1964). Interaction of arousal and recall interval in nonsense syllable paired-associate learning. Journal of Experimental Psychology, 67, 124-126.

KROLL, N.E.A., PARKS, T., PARKINSON, S.R., BIEBER, SL. and JOHNSON, A.L. (1970). Short-term memory while shadowing: recall of visually and aurally presented letters. Journal of Experimental Psychology, 85, 220-224.

LEA, S.E.G. and MORGAN, M. J. (1972). The measurement of rate-dependent changes in responding. In R.M. Gilbert and J.R. Millenson (Eds.), Reinforcement: Behavioural analyses (pp.129-146). London: Academic Press.

LETT, B.T. (1973). Delayed reward learning: Disproof of the traditional theory. Learning and Motivation, 4, 237-246.

LETT, B.T. (1974). Visual discrimination learning with a 1-min delay of reward. Learning and Motivation, 5, 174-181.

LETT, B.T. (1975). Long delay learning in the T-maze. Learning and Motivation, 6, 80-90.

LETT, B.T. (1977). Regarding Robert's reported failure to obtain visual discrimination learning with delayed reward. Learning and Motivation, 8, 136-139.

LETT, B.T. (1978). Long-delay learning of a black-white discrimination: Effects of varying the length of delay. Bulletin of the Psychonomic Society, 12, 307-310.

LETT, B.T. (1979). Long-delay learning: Implications for learning and memory theory. In N.S. Sutherland (Ed.), Tutorial essays in psychology: A guide to recent advances (Vol.2). New York: Halstead Press.

LIEBERMAN, D.A., McIntosh, D.C. and THOMAS, G.V. (1979). Learning when reward is delayed: A marking hypothesis. Journal of Psychology: Animal Behaviour Processes, 5, 224-242.

LOESS, H. (1968). Short-term memory and item similarity. Journal of Verbal Learning and Verbal Behaviour, 7, 87-92.

LOESS, H. and WAUGH, N.C. (1967). Short-term memory and intertrial interval. Journal of Verbal Learning and Verbal Behaviour, 6, 455-460.

MACKINTOSH, N.J. (1974). The psychology of animal learning. London: Academic Press.

MACKINTOSH, N.J. (1983). Conditioning and associative learning. Oxford: Clarendon Press.

MAKI, W.S. (1979). Discrimination learning without short-term memory: Dissociation of memory processes in pigeons. Science, 204, 83-85.

MAKI, W.S. Jr., and LEITH, C.R. (1973). Shared attention in pigeons. Journal of the Experimental Analysis of Behaviour, 19, 345-349.

MAKI, W.S. Jr., and LEUIN, T.C. (1972). Information processing by pigeons. Science, 176, 535-536.

McGAUGH, J.L. (1973). Drug facilitation of learning and memory. Annual Review of Pharmacology, 13, 229-241.

MUENZINGER, K.F. (1934). Motivation in learning: I. Electric shock for correct responses in the visual discrimination habit. Journal of Comparative Psychology, 17, 267-277.

MUENZINGER, K.F. and NEWCOMB, H. (1936). Motivation in learning: V. The relative effectiveness of jumping a gap and crossing an electric grid in a visual discrimination habit. Journal of Comparative Psychology, 21, 95-104.

MUENZINGER, K.F. and WOOD, A. (1935). Motivation in learning: IV. The function of punishment as determined by its temporal relation to the act of choice in the visual discrimination habit. Journal of Comparative Psychology, 20, 95-106.

NEURINGER, A.J. and CHUNG, S.H. (1967). Quasi reinforcement: Control of responding by a percentage reinforcement schedule. Journal of the Experimental Analysis of Behaviour, 10, 45-54.

NEVIN, J.A. (1979). Overall matching versus momentary maximizing: Nevin (1969) revisited. Journal of Experimental Psychology: Animal Behaviour Processes, 5, 300-306.

PALMERINO, C.C., RUSINIAK, K.W. and GARCIA, J.G. (1980). Flavour-illness aversions: The peculiar roles of odor and taste in memory for poison. Science, 208, 753-755.

PARE, W. (1961). The effect of caffeine and secondal on a visual discrimination task. Journal of Comparative and Physiological Psychology, 54, 506-509.

PEARCE, J.M. and HALL, G. (1978). Overshadowing the instrumental conditioning of a lever-press response by a more valid predictor of the reinforcer. Journal of Experimental Psychology: Animal Behaviour Processes, 4, 356-367.

PEARCE, J.M., NICHOLAS, D.J. and DICKINSON, A. (1981). The potentiation effect during serial conditioning. Quarterly Journal of Experimental Psychology, 33B, 159-179.

PETERSON, L.R. and GENTILE, A. (1965). Proactive interference as a function of time between tests. Journal of Experimental Psychology, 70, 473-478.

PETERSON, L.R., RAWLINGS, L. and COHEN, C. (1977). The internal construction of spatial patterns. In G.H. Bower (Ed.), The psychology of learning and motivation, vol. 11 (pp.245-276). New York: Academic Press.

RESCORLA, R.A. (1982). Effect of a stimulus intervening between CS and US in autoshaping. Journal of Experimental Psychology: Animal Behaviour Processes, 8, 131-141.

REVUSKY, S. (1971). The role of interference in association over a delay. In W.K. Honig and P.H.R. James (Eds.), Animal memory. New York: Academic Press.

RICHARDS, R.W. (1981). A comparison of signaled and unsignaled delay of reinforcement. Journal of the Experimental Analysis of Behaviour, 35, 145-152.

ROBERTS, J.E., TARPY, R.M. and LEA, S.E.G. (1984). Stimulus-response overshadowing: Effects of signaled reward on instrumental responding as measured by response rate and resistance to change. Journal of Experimental Psychology: Animal Behaviour Processes, 10, 244-255.

ROBERTS, W.A. (1976). Failure to replicate visual discrimination learning with a 1-min delay of reward. Learning and Motivation, 7, 313-325.

ROBERTS, W.A. (1977). Still no evidence for visual discrimination: A reply to Lett. Learning and Motivation, 8, 140-144.

ROBERTS, W.A. (1980). Distribution of trials and intertrial retention in delayed matching to sample with pigeons. Journal of Experimental Psychology: Animal Behaviour Processes, 6, 217-237.

ROBERTS, W.A. and GRANT, D.S. (1974). Short-term memory in the pigeon with presentation time precisely controlled. Learning and Motivation, 5, 393-408.

ROBERTS, W.A. and GRANT, D.S. (1978). An analysis of light-induced retroactive inhibition in pigeon short-term memory. Journal of Experimental Psychology: Animal Behaviour Processes, 4, 219-236.

SALZBERG, P.M., PARKS, T.E., KROLL, N.E.A. and PARKINSON, S.R. (1971). Retroactive effects of phonemic similarity on short-term recall of visual and auditory stimuli. Journal of Experimental Psychology, 91, 41-46.

SCHADE, A.F. and BITTERMAN, M.E. (1966). Improvement in habit reversal as related to dimensional set. Journal of Comparative and Physiological Psychology, 62, 43-48.

SHIMP, C.P. (1969). Optimum behaviour in free-operant experiments. Psychological Review, 76, 97-112.

SIZEMORE, O.J. and LATTAL, K.A. (1978). Unsignalled delay of reinforcement in variable-interval schedules. Journal of the Experimental Analysis of Behaviour, 30, 169-175.

SKINNER, B.F. (1938). The behaviour of organisms. New York: Appleton-Century-Crofts.

SKINNER, B.F. (1948). Superstition in the pigeon. *Journal of Experimental Psychology*, *38*, 168-172.

SKINNER, B.F. (1950). Are theories of learning necessary? *Psychological Review*, *57*, 193-216.

SPEAR, N.E. (1973). Retrieval of memory in animals. *Psychological Review*, *80*, 163-194.

SPEAR, N.E. (1978). *The processing of memories: Forgetting and retention*. New York: Erlbaum.

ST. CLAIRE-SMITH, R. (1979a). The overshadowing of instrumental conditioning by a stimulus that predicts reinforcement better than the response. *Animal Learning and Behaviour*, *7*, 224-228.

ST. CLAIRE-SMITH, R. (1979b). The overshadowing and blocking of punishment. *Quarterly Journal of Experimental Psychology*, *31*, 51-61.

SUTHERLAND, N.S. and MACKINTOSH, N.J. (1971). *Mechanisms of animal discrimination learning*. London: Academic Press.

TARPY, R.M., LEA, S.E.G. and MIDGLEY, M. (1983). The role of response-US correlation in stimulus-response overshadowing. *Quarterly Journal of Experimental Psychology*, *35B*, 53-65.

TARPY, R.M., ROBERTS, J.E., LEA, S.E.G. and MIDGLEY, M. (1984). The stimulus-response overshadowing phenomenon with VI versus FI schedules of reinforcement. *Animal Learning and Behaviour*, *12*, 50-54.

TELEGDY, G.A. and COHEN, J.S. (1971). Cue utilization and drive level in albino rats. *Journal of Comparative and Physiological Psychology*, *75*, 248-253.

THOMAS, D.R. (1969). The use of operant conditioning techniques to investigate perceptual processes in animals. In R.M. Gilbert and N.S. Sutherland (Eds.), *Animal discrimination learning*. London: Academic Press.

THOMAS, G.V., LIEBERMAN, D.A., McINTOSH, D.C. and RONALDSON, P. (1983). The role of marking when reward is delayed. *Journal of Experimental Psychology: Animal Behaviour Processes*, *9*, 401-411.

TOLMAN, E.C. (1948). Cognitive maps in rats and men. *Psychological Review*, *55*, 189-208.

TRANBERG, D.K. and RILLING, M. (1980). Delay-interval illumination changes interfere with pigeon short-term memory. *Journal of the Experimental Analysis of Behaviour*, *33*, 39-49.

TULVING, E. and OSLER, S. (1968). Effectiveness of retrieval cues in memory for words. *Journal of Experimental Psychology*, *77*, 593-601.

WAGNER, A.R. (1969). Incidental stimuli and discrimination learning. In R.M. Gilbert and N.S. Sutherland (Eds.), Animal discrimination learning (pp. 83-111). London: Academic Press.

WAGNER, A.R. (1976). Priming in STM: An information-processing mechanism for self-generated or retrieval-generated depression in performance. In T.J. Tighe and R.N. Leaton (Eds.), Habituation: Perspectives from child development, animal behaviour, and neurophysiology (pp.95-128). Hillsdale, NJ: Erlbaum.

WAGNER, A.R. (1978). Expectancies and the priming of STM. In S.H. Hulse, H. Fowler and W.K. Honig (Eds.), Cognitive processes in animal behaviour. Hilldale, NJ: Erlbaum.

WAGNER, A.R. (1979). Habituation and memory. In A. Dickinson and R.A. Boakes (Eds.), Mechanisms of learning and motivation: A memorial to Jerzy Konorski. Hillsdale, NJ: Erlbaum.

WAGNER, A.R. (1981). SOP: A model of automatic memory processing in animal behaviour. In N.E. Spear and R.R. Miller (Eds.), Information processing in animals: Memory mechanisms. Hillsdale, NJ: Erlbaum.

WAGNER, A.R., RUDY, J.W. and WHITLOW, J.W. (1973). Rehearsal in animal conditioning. Journal of Experimental Psychology, 97, 407-426. (Monograph).

WALKER, E.L. and TARTE, R.D. (1963). Memory storage as a function of arousal and time with homogenous and heterogenous lists. Journal of Verbal Learning and Verbal Behaviour, 2, 113-119.

WERNER, H. and WAPNER, S. (1952). Towards a general theory of perception. Psychological Review, 59, 324-338.

WICKELGREN, W.A. and NORMAN, D.A. (1966). Strength models and serial position in short-term recognition memory. Journal of Mathematical Psychology, 3, 316-347.

WICKENS, D.D., BORN, D.G. and ALLEN, C.K. (1963). Proactive inhibition and item similarity in short-term memory. Journal of Verbal Learning and Verbal Behaviour, 2, 440-445.

WILLIAMS, B.A. (1976). The effects of unsignalled delayed reinforcement. Journal of the Experimental Analysis of Behaviour, 26, 441-449.

WINER, B.J. (1971). Statistical principles in experimental design (2nd ed.). New York: McGraw-Hill.

WORSHAM, R.W. and D'AMATO, M.R. (1973). Ambient light, white noise, and monkey vocalisation as sources of interference in visual short-term memory of monkeys. Journal of Experimental Psychology, 99, 99-105.

YOUNG, D.B. (1979). Visual discrimination learning with delayed reward, when the choice response is 'marked'. Unpublished B.Sc. dissertation, Stirling University.

ZEILER, M. (1977). Schedules of reinforcement: The controlling variables. In W.K. Honig and J.E.R. Staddon (Eds.), *Handbook of operant behaviour* (pp.201-232). Englewood Cliffs, N.J.: Prentice-Hall.